DDoS Attack Mitigation

FortiDDoS™ 4.1

Handbook

FortiDDoS 4.1 Handbook

June 10, 2014

1st Edition

Technical Documentation	http://docs.fortinet.com
Knowledge Base	http://kb.fortinet.com
Forums	https://support.fortinet.com/forum
Customer Service & Support	https://support.fortinet.com
Training	http://training.fortinet.com
FortiGuard Threat Research & Response	http://www.fortiguard.com
License	http://www.fortinet.com/doc/legal/EULA.pdf
Document Feedback	Email: techdocs@fortinet.com

Table of contents

Introduction

FortiDDoS is a network behavior anomaly (NBA) prevention system that detects and blocks network attacks that are characterized by excessive use of network resources, commonly referred to as distributed denial of service (DDoS) attacks.

FortiDDoS uses a variety of schemes, including anomaly detection and statistical techniques, to provide nonstop protection, even against attacks cannot be recognized by an attack signature yet. It focuses on intent rather than content of network attacks. When it detects an intrusion, FortiDDoS immediately blocks traffic, thus protecting the systems it defends from being flooded.

Product features

The following features make FortiDDoS the best in its class.

- Initial learning period

  Your FortiDDoS learns based on traffic patterns to and from one or more protected systems. It needs about 2 to 7 days of attack-free learning to collect enough information to protect your systems. During this period, the appliance operates in detection mode. In detection mode, FortiDDoS does not drop any packets and operates with high (factory default) thresholds. The length of the initial learning period depends upon the seasonality of traffic (its predictable or expected variations) and how representative of normal traffic conditions the learning period is.

  Weekends alone do not make good learning periods for businesses that have substantially different traffic during the week. Thus, it is better to start the learning period on a weekday. In most cases, 7 days is sufficient to capture the weekly seasonality in traffic.

  At the end of the initial learning period, you can lower the thresholds to the values that are recommended by FortiDDoS and, still in detection mode, monitor the anomalies that the appliance detects for false postives and false negatives. If needed, adjust thresholds and monitor the results.

  When you are satisfied with the appliance settings, put the system into prevention mode, in which the appliance blocks and drops traffic. For additional tuning options, see the information about Service Protection Profiles found in this handbook.
- Continuous learning

  The FortiDDoS begins learning traffic patterns as soon as it begins monitoring traffic, and it never stops learning. By analyzing traffic rates, FortiDDoS dynamically sets thresholds to differentiate between legitimate traffic increases and attacks.
- ‘Zero Day’ attack prevention

  Because the FortiDDoS uses rate-based analysis, it provides protection against attacks that hackers haven’t even thought up yet. Administrators do not need to intervene and the appliance is “on guard” 24/7, automatically protecting your network systems and bandwidth.
- Configurable event monitoring

  You can monitor FortiDDoS events using the web UI, SNMP, or email event notification.
- Viewing dropped packet statistics

FortiDDoS allows you to view graphs which display counts of the packets it has dropped. In detection mode, FortiDDoS allows you to monitor packets which it would drop if it were in prevention mode.

- Purpose-built for low latency and rapid response

  FortiDDoS's patented combination of purpose-built hardware and heuristics allows the appliance to function inline (for example, between the external network and a protected server), where it can operate at a high rate to detect and mitigate attacks, even when an attack is underway. FortiDDoS introduces a latency of approximately a few microseconds and has a response time of 2 seconds or less.

- Intuitive analysis tools and reports

  The built-in reporting tools show graphical analysis of network traffic history from five minutes to one year. You can view traffic profiles using a broad range of Layer 3, 4 or 7 parameters. Event reporting tools show tabular and graphical analysis of attack events for a range of time periods, from the past hour to the past year. With just a few clicks, you can create intuitive and useful reports such as top attackers, top attacks, top attack destinations, top connections, and so on.

- Granular monitoring and configuration

  FortiDDoS's custom hardware design monitors thresholds for all traffic it sees on Layers 3, 4, and 7. It measures byte and packet counts, state transitions, fragments, checksums, flags, new connections, address pairs, and so on. You can set thresholds on any traffic parameter to limit traffic for particular systems or applications.

- Source tracking

  In addition to blocking attacks, FortiDDoS pinpoints and logs the originators of attacks for further administrative action.

- Aggressive aging under attack

  Many botnets have started using slow connection build up as a mechanism to confuse security appliances and thus effectively overload the servers. FortiDDoS can identify these types of attacks by maintaining concurrent connections per source and concurrent connections per destination thresholds. When these limits are exceeded, FortiDDoS aggressively ages the targeted connections and thus reduce the load.

- IP address matching

  A proprietary algorithm matches incoming connection requests with known IP addresses to mitigate SYN attacks without the overhead of connection proxy. Legitimate users can connect or remain connected, even during a SYN attack.

- Service Protection Profiles (SPPs)

  FortiDDoS can maintain up to 8 sets of counters and thresholds that you assign to a subnet as a group. Thus, a single appliance can protect up to 8 subnets, each identified by an IP address representing a server or group of networked servers. Each of these virtual protection zones - called Service Protection Profiles -- learns traffic patterns and estimates adaptive thresholds independently. You can assign each profile an independent administrator, which is useful in environments such as an Internet service provider.

- Deep packet inspection

  The deep packet inspection provided by the FortiDDoS hardware allows it to identify certain header fields in HTTP packets. Using this capability, FortiDDoS can identify DDoS attacks on specific URLs and prevent other URLs in the system from getting overwhelmed due to attacks on just one or a few locations.

Architecture

FortiDDoS protects a system or a network of systems from rate-based and anomaly attacks. In the simplest configuration, you can install a FortiDDoS appliance as an inline device, as shown in Figure 1. It monitors both inbound traffic arriving from the external network (usually the Internet) and outbound traffic from the protected system (for example, a server or LAN).

Figure 1: Network with a FortiDDoS appliance protecting a single system

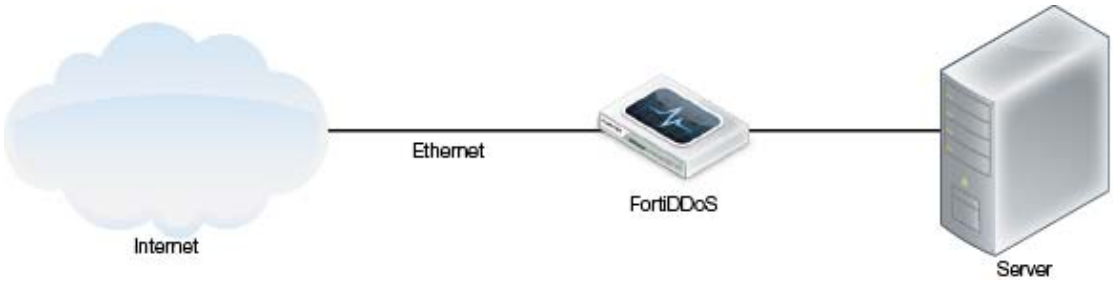

In addition to this serial or inline arrangement, you can deploy FortiDDoS in more complex and specialized topologies. See “Planning the network topology” on page 54.

Scope

This document describes how to set up your FortiDDoS appliance. It describes how to complete first-time system deployment, including planning the network topology.

It also contains the instructions on how to:

- Use the web user interface (web UI).
- Specifying the IP addresses and address ranges to protect by creating System Protection Profiles (SPPs).
- Work with system-generated statistics to configure and adjust the attack protection settings for each profile.

After completing “How to set up your FortiDDoS” on page 53:

- You will have administrative access to the web UI, CLI, or both.
- You will have completed firmware updates, if any.
- The system time, DNS settings, administrator password, and network interfaces will be configured.
- You will have created a Service Protection Profile (SPP).
- You will have specified an initial operating mode.
- You will have completed an initial configuration for the attack protection thresholds.

Once that basic installation is complete, you can use the rest of this document to use the web UI to:

- Monitor attack activity and create reports.
- Adjust thresholds to optimize traffic flow and eliminate false positives.
- Customize the web UI.
- Diagnose problems.

What's new

The following features are new or changed since FortiDDoS 3.2:

FortiDDoS 4.1

- **Logging & report enhancements**
  - **SNMP traps & MIBs for attack logs** — You can now configure FortiDDoS to send attack log information to SNMP managers. See "Configuring SNMP settings for attack log messages" on page 248.
  - **DDoS Subnet Attack Log** — The new *DDoS Subnet Attack Log* displays events associated with a specific SPP policy, with counts updated every five minutes. See "DDoS Attack Log and DDoS Subnet Attack Log" on page 222.
  - **Subnet Executive Summary dashboard** — The new *Subnet Executive Summary* dashboard displays all attacks in the *Top Attacked Subnet* and *Top ACL Subnet Drops* report categories. See "Viewing report information on a dashboard (Executive Summary)" on page 250.
  - **Destination tracking** — For all attack log event categories, FortiDDoS now provides the IP address of the first destination it identifies as the target of the attack activity. Information that is organized by this destination is available as a report type and widgets on the *Executive Summary* and *Attack Graphs* dashboards.
  - **Filter report information by SPP or subnet** — When you create a report configuration, you can now restrict the information in the report to a specific service protection profile (SPP) or subnet. See "DDoS Attack Activity report types" on page 255.
- **Enhanced blocking by geolocation** — The *Geo Location Policy* setting allows you to either permit traffic from all geographic locations and add exceptions or deny access to all locations with exceptions. See "Blocking addresses from a specific geographic location, anonymous proxies, and satellite providers" on page 120.
- **Access dropped packet and other statistics via API** — You can now use the FortiDDoS REST API to access dropped and blocked traffic statistics and traffic graph information. See the *FortiDDoS REST API Reference*.
- **MySQL access to DDoS attack log** — You can now access the DDoS attack log with read-only permission using a third-party tool such as the MySQL command-line tool or MySQL Workbench. See "Accessing the DDoS attack log using SQL" on page 231.
- **Alert email message for SPP switching** — You can now configure FortiDDoS to generate a system event log and send a corresponding email message whenever the appliance switches a subnet to its alternative service protection profile (SPP). (If you want FortiDDoS to notify you that the traffic level has exceeded the SPP switching threshold without switching the SPP, in the SPP policy settings, specify the same SPP for both *Service Protection Profile* and *Alternate Service Protection Profile*.) See "Selecting which system events to log" on page 236 and "Alert email" on page 241.
- **Improved dual-stack IPv6 support** — Additional settings and functionality that make it easier to deploy FortiDDoS in networks with IPv6 traffic. See "Enabling Internet Protocol version 6 (IPv6) support" on page 107.
- **Double VLAN (DVLAN) detection** — FortiDDoS now tracks traffic with an additional 802.1Q tag (for example, VLAN Q-in-Q).

FortiDDoS 4.0.1

- No design changes. Bug fixes only.

FortiDDoS 4.0

- **Additional data ports** — 16 physical LAN and WAN ports are configured as linked pairs. Odd-numbered ports are LAN connections that have a corresponding even-numbered port, which is the associated WAN connection. That is, Port 1/Port 2 behaves as LAN 1/WAN 1, Port 3/Port 4 as LAN 2/WAN 2, Port 5/Port 6 as LAN 3/WAN 3, and so on. These port pairs enable you to protect up to 8 links with a single appliance.
- **Increased throughput** — Each WAN/LAN link pair has a maximum throughput of 1 Gbps full duplex. Each model has the following total appliance throughput:
  - **FortiDDoS 400B:** 4 Gbps full duplex
  - **FortiDDoS 800B:** 8 Gbps full duplex
  - **FortiDDoS 1000B:** 12 Gbps full duplex
  - **FortiDDoS 2000B:** 24 Gbps full duplex
- **Configuration synchronization** — High Availability (HA) configuration allows you to synchronize configuration information between two FortiDDoS appliances to create a secondary appliance that always has an up-to-date configuration.

  See "Topology for synchronizing the configuration of two FortiDDoS appliances" on page 64.
- **Automatic bypass for copper links** — For Ethernet links (copper, RJ-45), the FortiDDoS appliance automatically passes traffic through when the appliance is not powered up, its FortiASIC processor or integrated switch fabric fail, or it is booting up and all services are not yet available.

  See "External bypass switches for maintenance & failover" on page 55.
- **Link down synchronization** — The appliance has two options for Link Down Synchronization: Wire and Hub. When Wire is selected, FortiDDoS monitors the link state of both ports in a port pair. If the link goes down on either port, it disables the other port. The appliance re-enables the port when it detects that the link for other port in the pair is up again. When Hub is selected, FortiDDoS does not disable both ports in a port pair if the link goes down on one of the ports.

  See "Configuring bypass mode" on page 186.
- **Redesigned web UI—** The graphical user interface is organized by component and tasks. Many of its system settings and options are shared with other Fortinet products.
- **Management via command-line interface** — You can perform all appliance configuration from a Secure Shell (SSH) or Telnet terminal or from the JavaScript CLI Console widget in the web UI.

  See "Connecting to the web UI or CLI" on page 74.
- **RESTful web API configuration** — Use a web API that uses HTTP and REST principles to perform tasks such as allowing or denying sources, setting thresholds and changing Service Protection Profile (formerly VIDs) configuration.

  For more information, contact Fortinet Technical Support.
- **BIOS-based signed appliance certificate** — The validation mechanism for the appliance's identity is built into its hardware.

  See "Certificate configuration" on page 187.
- **Faster threshold report generation** — FortiDDoS now takes less time to generate the traffic statistics it uses to calculate system-recommended thresholds.

  See "Generating and reviewing a traffic statistics report" on page 144.
- **Save reports as PDF** — Many FortiDDoS system events and attack activity reports and graphs have a *Save as PDF* option that exports information in a format that is suitable for printing and sharing.
- **Attack activity at a glance dashboard** — Access the most popular attack activity reports information on a single web page and in table format using *Log & Report* > *Report Browse* >

Executive Summary. Use *Log & Report* > *Attack Graphs* > *Attack Graphs* to access the most popular attack activity graphs on a single web page.

- **Context-sensitive help** — Click *Help* to open the HTML help information for the current content pane.
- **Filter and sort log information** — For system event and DDoS attack logs, you can use the column headers to sort log information or arrange the columns. The filter feature allows you to select items to include or exclude based on date, category, or other criteria.

  See "Viewing log messages" on page 239.
- **Enhanced reports** — New features include the ability to generate reports as HTML, text, PDF and customize reports with a logo.

  See "Customizing the report's headers, footers, & logo" on page 252.
- **Alternate image support** — You can now install alternate FortiDDoS firmware loaded from a separate partition on the appliance's hard disk.

  See "Installing alternate firmware" on page 87.
- **Built-in DoS control** — FortiDDoS blocks packets with a pre-defined set of anomalies before they reach the appliance's processor. Traffic graphs and reports do not report the packets that this feature drops.

  See "Installing alternate firmware" on page 87.
- **Block protocols on subnets (Distress ACL)** — The distress ACL feature helps to block brute force protocol attacks on a specified subnet or IP address. It allows you to block packets that can flood the pipe before they reach the appliance's processor. Traffic graphs and reports do not report the packets that this feature drops.

  See "Preset access control" on page 118.
- **Bypass fiber ports for 2000B model** — Two physical port pairs on the FortiDDoS 2000B have built-in bypass capability. Built-in bypass works during a power failure, critical component failure and during startup and shutdown.

  For more information, see *FortiDDoS 2000B QuickStart Guide*.
- **IP Reputation update using file upload** — You can update the addresses in the IP Reputation Service list by uploading a .pkg file.

  See "License Information widget" on page 198.

Documentation enhancements

Information from the *Fundamentals Guide*, *Install Guide*, and *Web-based Manager Reference Guide* are consolidated in this handbook.

Key concepts

This chapter defines basic FortiDDoS concepts, terms, and features.

If you are new to FortiDDoS, or new to distributed denial of service (DDoS), this chapter can help you to quickly understand.

See also

- Planning the network topology

Network behavior analysis (NBA)

Computer network security is a challenge as old as the Internet itself. The sophistication and infamy of network-based system attacks has kept pace with the security technology and hackers only feel more challenged by the latest heuristics designed to foil their efforts.

Some attackers exploit system weaknesses for political purposes, disgruntled about the state of software or hardware in the market today. Others target specific systems out of spite or a grudge against a specific company.

Yet others are simply in search of the infamy of bringing a high-traffic site to its knees with a denial of service (DoS) attack. In such an attack, the hacker attempts to consume all the resources of a networked system so that no other users can be served. The implications for victims range from a nuisance to millions of dollars in lost revenue.

Consequences of attacks

Any computer can be infected, and the consequences can range from a nuisance popup ad to thousands of dollars in costs for replacement or repair. For this reason, AntiVirus (AV) software for all PCs should be a mandatory element of any network security strategy. But whether you measure cost in terms of lost revenue, lost productivity, or actual repair/restore expenses, the cost of losing a server to an attack is far more severe than losing a laptop or desktop.

Servers that host hundreds or thousands of internal users, partners, and revenue-bearing services are usually the targets of hackers, because this is where the pain is felt most. Protecting these valuable assets appropriately is paramount. In early 2000, the industry saw a new kind of 'worm' attack, in which hundreds or thousands of (sometimes unsuspecting) systems were employed to simultaneously bombard a target host, paralyzing its productivity. Several high traffic sites such as Amazon.com, Buy.com, CNN, Yahoo, and eBAY were affected by these Distributed Denial of Service (DDoS) attacks.

Because each attacking system looks innocent, advanced techniques are required to separate the 'bad' traffic from the 'good' traffic.

Distributed denial of service (DDoS) attacks

In distributed denial of service (DDoS) attacks, hackers write a program that will covertly send itself to dozens, hundreds, or even thousands of other computers. These computers are known as 'agents' or 'zombies', because they act on behalf of the hackers to launch an attack against target systems. A network of these computers is called a botnet.

To circumvent detection, attackers are increasingly mimicking the behavior of a large number of clients. The resulting attacks are hard to defend against, using standard techniques, as the malicious requests differ from the legitimate ones in intent but not in content.

At a predetermined time, the worm will cause all of these zombies to attempt repeated connections to a target site. If the attack is successful, it will deplete all system or network resources, thereby denying service to legitimate users or customers.

E-commerce sites, domain name servers, web servers, and email servers are all vulnerable to these types of attacks. IT managers must take steps to protect their systems - and their businesses- from irreparable damage.

Strategies for protection

The best security strategies encompass people, operations, and technology. The first two typically fall within an autonomous domain, e.g. within a company or IT department that can enforce procedures among employees, contractors or partners. But since the Internet is a public resource, such policies cannot be applied to all potential users of a public website or email server. Thankfully, technology offers a range of security products to address the various vulnerabilities.

Firewalls

Firewalls can go a long way to solving some problems by restricting access to authorized users and blocking unwanted protocols. As such, they are a valuable part of a security strategy. But public web sites and eCommerce servers cannot know in advance who will access them and cannot 'prescreen' users via an access list. Certain protocols can be blocked by firewalls, but most DoS attacks utilize authorized ports (e.g TCP port 80 for a web server) that cannot be blocked by a firewall without effectively blocking all legitimate HTTP traffic to the site, thereby completing the hacker's task.

Firewalls offer some security against a single user DoS attack by denying access to the offending connection (once it is known), but most DoS attacks today are distributed among hundreds or thousands of zombies, each of which could be sending legal packets that would pass firewall scrutiny. Firewalls perform a valuable service in an integrated security strategy, but firewalls alone are not enough.

Router access control lists

Likewise, access lists in the router can be used to block certain addresses, if such addresses can be known a priori. But web sites open to the public are, by nature, open to connections from individual computers, which are exactly the agents hackers use to initiate attacks. In a Distributed DoS (DDoS) Attack, thousands of innocent looking connections are used in parallel. Although router access lists can be used to eliminate offending packets once they are identified, routers lack the processing power and profiling heuristics to make such identifications on their own.

In addition, complex access lists can cause processing bottlenecks in routers, whose main function is to route IP packets. Performing packet inspections at layers 3, 4, and 7 taxes the resources of the router and can limit network throughput.

Antivirus software

End systems cannot be considered secure without AntiVirus software. Such software will scan all inputs to the system for known viruses and worms, which can cause damage to the end system and any others they may infect. Even after a virus is known and characterized, instances

of it are still circulating on the Internet, through email, on compact discs (CDs) and floppy disks. A good AntiVirus subscription that is frequently updated for the latest protection is invaluable to any corporate or individual computer user.

But even AntiVirus software is not enough to catch certain attacks that have been cleverly disguised. Once a system is infected with a new strain, the damage can be done before the virus or worm is detected and the system is disinfected.

Application protection

Such packages include software that watches for email anomalies, database access queries, or other behavior that may exploit vulnerability in the application. Because it must be very specific - and very close - to the application it is protecting, application protection is typically implemented as software on the host. Dedicated servers would benefit from well-designed application security software that will maintain the integrity of the code and detect anomalous behavior that could indicate an attack. Certain malicious code can attempt to overwrite registers on the end-system and thereby hijack the hardware for destructive purposes.

Intrusion detection systems

Intrusion Detection Systems (IDS) are designed to 'listen' to traffic and behavior and set an alarm if certain conditions are met. Some IDS implementations live in the host, while others are deployed in the network. The IDS sensor monitors traffic, looking for protocol violations, traffic rate changes or matches to known attack 'signatures'. When a threat is detected, an alarm is sent to notify a (human) network administrator to intervene.

All IDSs use software, but some run on general purpose computers, while others make use of purpose-built hardware.

Host-based intrusion detection and prevention systems

Some intrusion detection systems are designed as software running on general purpose computing platforms. Not to be confused with application security software (mentioned above), which runs on the end system and focuses primarily on layer 5-7, software based intrusion systems must also focus on layers 3 and 4 of the protocol stack. These packages rely on the CPU power of the host system to analyze traffic as it comes into the server. General purpose computers often lack the performance required to monitor real-time network traffic and perform their primary functions. Creating a bottleneck in the network or on the server actually helps the hacker accomplish his goal by restricting access to valuable resources.

End-systems provide the best environment for signature recognition because packets are fully reassembled and any necessary decryption has been performed. However, signature based intrusion detection has its limitations, as described below.

Content-based intrusion prevention systems

The next step in the evolution of Intrusion security leads to Intrusion Prevention Systems (IPS). Unlike Intrusion Detection Systems, which require manual intervention from an administrator to stop an attack, an IPS will automatically take action to prevent an attack once it is recognized. This can cut down response time to near zero, which is the ultimate goal of intrusion security.

Intrusion Prevention must be intelligent, however, or the remedy may actually accomplish the hacker's goal of denying resource to legitimate users.

Prevention mechanisms can also be harmful if detection is subject to false positives, or incorrect identification of intrusion. If the prevention action is to disable a port, protocol, or address, a false positive could result in denial of service to one or more legitimate users.

Network behavior analysis

An alternative to signature recognition is network behavior analysis. Rate-based systems must provide detailed analysis and/or control of traffic flow. A baseline of traffic patterns is established, usually during a learning mode in which the device only 'listens' without acting on any alarm conditions. A good system will have default parameters set to reasonable levels, but the 'listening' period is required to learn the traffic behavior on various systems. The listening period should be 'typical,' in the sense that no attacks or unusual traffic patterns should be present. For example, Saturday and Sunday are probably not good days to build a baseline for a corporate server that is much busier during the workweek. Periods of unusually high or low traffic also make bad listening intervals, such as Christmas vacation week, unusually high traffic due to external events (press releases, sales promotions, Super Bowl halftime shows, and so on).

Once a baseline is established, rate-based systems watch for deviations from the known traffic patterns to detect anomalies. Good systems will allow an administrator to override the baseline parameters if events causing traffic surges are foreseen, for example, a server backup scheduled overnight.

While signature-based systems are scrutinized for false-negatives, or failing to identify an attack, rate-based systems should be scrutinized for false positives, or misidentifying legitimate changes in traffic patterns as attacks. Whether setting alarms or taking preventative action, rate-based systems must be well-designed to avoid unnecessary overhead.

Equally important for rate-based systems are their analysis tools. Administrators should be able to view their traffic patterns on a variety of levels, and use this information to tune their network resources.

Differences & similarities with conventional firewalls

Conventional stateful firewalls drop packets or stateful connections, but they cannot correlate packets to a source. FortiDDoS has a unique feature that allows it to promptly correlate attacks and verify if they are initiated by a single host. If it can do that (in case it is a non-spoofed attack), it blocks the offending source for a longer period of time.

It is important to understand the differences between a stateful firewall and a stateful Network Behavior Analysis (NBA) system such as FortiDDoS. Here are the key differences: Conventional stateful firewalls have rules that allow or deny packets or individual connections based on their individual characteristics. They do not remember packets in an aggregate way.

FortiDDoS operates on an aggregate basis. It looks at packet rates – typically within one second, over a period of time. It measures packet rates for various layer 3, 4, and 7 parameters and compares against thresholds set for them. If the rate exceeds the threshold, it blocks them for a configured period.

In a firewall, the administrator can set a rule which allows the UDP destination port 1434 regardless of the rate.

A FortiDDoS administrator, on the other hand, can set a rule which allows UDP 1434 only if the rate is within 10 packets per second. Beyond this rate, the UDP packets destined to that port are dropped.

There are some features in FortiDDoS which are similar to a firewall. It is therefore important to learn how to graduate to FortiDDoS.

Like a firewall, FortiDDoS allows you to configure layer 3, 4, and 7 blocking conditions. For a description of the traffic parameters that you can configure, see “Access control lists (ACLs)” on page 119.

Comparing FortiDDoS to conventional intrusion prevention systems (IPS)

FortiDDoS is a rate-based intrusion prevention system (IPS) device that detects and blocks network attacks which are characterized by excessive use of network resources. It uses a variety of schemes, including anomaly detection and statistical techniques, to detect and block malicious network traffic. When it detects an intrusion, the FortiDDoS blocks traffic immediately, thus protecting the systems it is defending from being overwhelmed.

Unlike conventional content-based intrusion prevention systems (CBIPS), a network behavior analysis (NBA) system does not rely on a predefined attack “signature” to recognize malicious traffic. A CBIPS is vulnerable to “zero-day” attacks, or attacks that cannot be recognized because no signature has been identified to match the attack traffic. In addition, attack traffic that is compressed, encrypted, or effectively fragmented can escape many pattern-matching algorithms in content-based systems. And many rate-based attacks are based on genuine and compliant traffic being sent at high rates, which effectively cheats the CBIPS.

An NBA provides a network with unique protection capabilities. It delivers security services not available from traditional firewalls, intrusion detection systems (network or host-based), or antivirus/spam detectors. The detection, prevention, and reporting of network attacks is based on traffic patterns rather than individual transaction or packet-based detection, which enables the FortiDDoS to serve a vital role in an effective security infrastructure. Rather than replacing these elements, an NBA complements their presence to form a defense-in-depth network security architecture.

Comparing FortiDDoS to conventional network behavior analysis (NBA)

Preventing DDoS attacks requires maintaining highly granular statistics at Layers 3, 4, and 7. It is essential to track individual sources, destinations, protocols, connections and ports that can add up to millions of parameters. FortiDDoS is a hardware-based NBA solution and unlike software-based solutions, it maintains normal levels of processing and data throughput during denial of service attacks.

FortiDDoS's custom hardware design monitors thresholds for all traffic on Layers 3, 4, and 7. It measures byte and packet counts, state transitions, fragments, checksum, flags, new connections, address pairs, and so on. You can set thresholds on any of these network parameters to limit the traffic rate for particular systems or applications.

To recognize and prevent attacks, FortiDDoS monitors dozens of parameters to detect subtle changes in the behavior of network traffic. The following table shows the capacity of counting mechanisms provided by FortiDDoS.

Table 1: FortiDDoS Traffic Threshold Table

Layer	Type	Counters
3	1. Protocol flood	256
	2. Fragment flood	1
	3. IP source flood & source tracking	1 million
	4. IP destination flood	1 million

Table 1: FortiDDoS Traffic Threshold Table

Layer	Type	Counters
4	1. TCP port flood	65535
	2. UDP port flood	65535
	3. ICMP type/code flood	65535
	4. TCP connection flood	1 million
	5. Legitimate IP table (for SYN and zombie flood)	2 million
	6. SYN flood	1
	7. Excessive SYN rate/source	1 million sources
	8. Excessive concurrent connections/source	1 million sources
	9. Excessive concurrent connections/destination	1
	10. Excessive ACK/destination	1
	11. Excessive RST/destination	1
	12. Excessive FIN/destination	1
7	1. HTTP methods	8
	2. URL Flood	65536
	3. Host	512
	4. Referer	1
	5. Cookie	1
	6.User-Agent	1
	7. Mandatory HTTP Headers	1
	8. Sequential Accesses	1
	9. URLs/source	1 million sources
	10. Excessive SIP INVITE requests/source	1 million sources
	11. Excessive concurrent SIP INVITE requests/source	1 million sources
	12. Excessive SIP REGISTER requests/source	1 million sources

Configuration workflow

Most of your work with FortiDDoS can be included in two main tasks: configuration and identifying and analyzing attacks.

For the information on how to respond when you detect an attack, see “Attack analysis workflow” on page 33.

Figure 2: Typical workflow for configuration

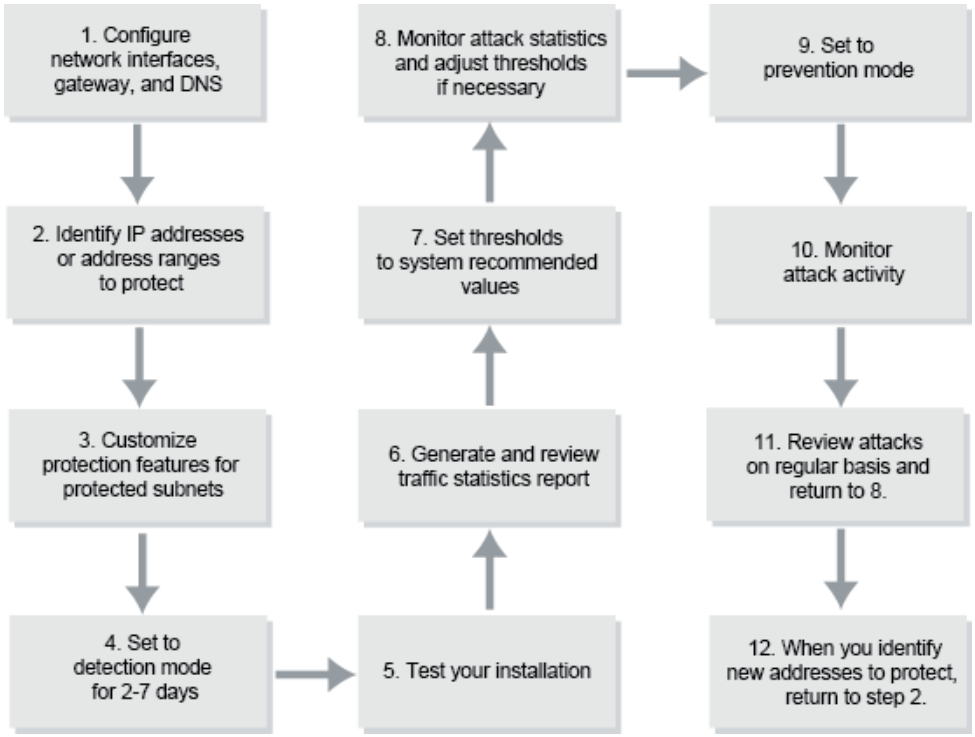

Detailed procedures for these steps are included in the following sections:

- “Configuring network interfaces, gateway, and DNS” on page 98
- “Identifying IP addresses and subnets to protect (SPP creation)” on page 111
- “Customizing protection features for protected subnets” on page 118
- “Setting FortiDDoS to detection mode” on page 116
- “Testing your installation” on page 141
- “Generating and reviewing a traffic statistics report” on page 144
- “Setting thresholds to system recommended values” on page 146
- “Monitoring attack statistics” on page 150 and “Adjusting thresholds” on page 150
- “Setting FortiDDoS to detection mode” on page 116 (set profile to prevention mode)
- “Monitoring attack activity and other system information” on page 195

See also

- How to set up your FortiDDoS

Anomalies that FortiDDoS blocks

FortiDDoS blocks the following three types of anomalies:

- **Header anomalies** — Header anomalies are Layer 3, 4, or 7 HTTP headers with content that violates the protocol's standards. For more information, see "Anomaly Drops graphs" on page 218.
- **TCP state anomalies** — State anomalies are found within the packets of a single TCP connection. For more information, see "TCP State Anomalies drop graph" on page 218.
- **Rate anomalies** — Rate anomalies are the main focus of FortiDDoS and represent the bulk of its functionality. Through continuous learning and adaptive threshold estimation, FortiDDoS set thresholds on rates for various Layer 3, 4, and 7 parameters in each direction. If these thresholds are violated, FortiDDoS considers the source traffic a rate anomaly and blocks it.

Other sections of this handbook discuss various rate anomalies that FortiDDoS detects and prevents in detail, including the following sections:

- "Continuous learning & adaptive threshold estimation" on page 23 (describes how FortiDDoS analyzes traffic and generates thresholds)
- "Adjusting thresholds" on page 150 (describes setting multiple and individual thresholds that limit traffic by rate)

Continuous learning & adaptive threshold estimation

Unlike other network behavior analysis (NBA) systems, FortiDDoS never stops learning. FortiDDoS continuously learns traffic patterns for a large group of layer 3, 4, and 7 parameters in both directions. This learning process ensures that the appliance takes into account the traffic's average, trend and seasonality (expected or predictable variations) at any moment before deciding what threshold to set for the parameter.

Thresholds for traffic

FortiDDoS sets traffic thresholds using specific parameters. For example, you can set a threshold on how many SYN packets reach a server per second before FortiDDoS identifies this traffic as a SYN flood. You can also set a threshold for fragmented packets and other parameters.

Traffic prediction

FortiDDoS uses the following information to predict traffic patterns:

- The historical base, or weighted average of recent traffic (FortiDDoS gives the most recent traffic more weight)
- The trend, or slope, of the traffic
- The seasonality of traffic over historical time periods (predictable or expected variations)

Figure 3: Trend, slope, and base of traffic

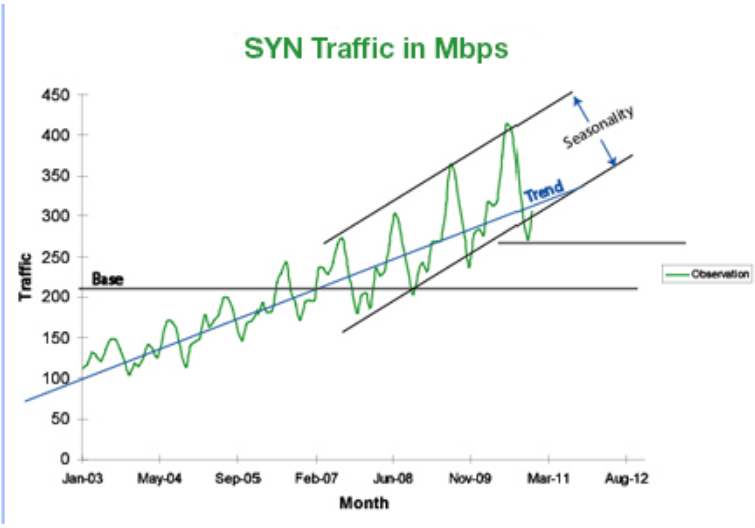

FortiDDoS uses these statistics to create a forecast for the next traffic period.

Figure 4: Forecast vs. actual traffic

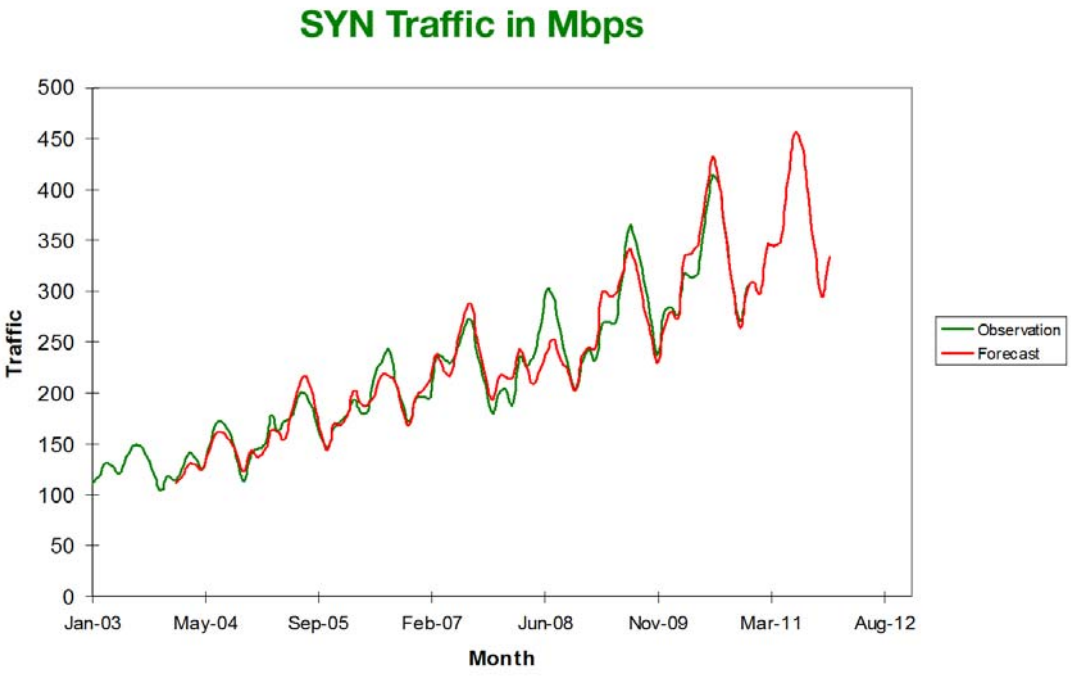

Because traffic is nondeterministic, the forecast cannot be exact. The extent to which an observed traffic pattern is allowed to exceed its forecast is bounded by a **threshold**.

Generally speaking, a threshold is a rate against which FortiDDoS compares observed traffic to determine whether a rate anomaly is occurring. This definition describes the thresholds

displayed by any of the traffic statistics. But the actual value for a threshold can vary depending on traffic rates and conditions. At any given time, the threshold may be one of the following:

- Configured minimum threshold
- Estimated threshold
- Adaptive limit maximum threshold
- A threshold adjusted using *Protection Profiles* > *System Recommendation*

Fixed vs. adaptive thresholds

Most NBA systems allow you to set thresholds as a fixed value. But most network administrators realize that traffic is never static. It shows trends and seasonality (a predictable or expected variation). FortiDDoS can apply thresholds adaptively — it has the intelligence to make adjustments based on past history.

FortiDDoS determines how many packets or frames is too many by defining limits, or thresholds, for each of the measurable parameters of network traffic. When traffic exceeds one or more thresholds, FortiDDoS makes some intelligent decisions about what actions to take.

FortiDDoS defines the thresholds for each of the monitored parameters associated with layers 3, 4, and 7.

You can also define Access Control Lists that block protocols, ports, addresses, or packets with specified header contents.

Configured minimum thresholds

Rate anomalies (that is, floods) are triggered by any traffic that exceeds a configured minimum threshold, which is a value that FortiDDoS provides for each threshold after the initial learning period. You can accept these values or configure FortiDDoS to use one of the following values as the minimum threshold:

- The configured minimum threshold adjusted by an amount that you specify in percent (using *System Recommendation*; see "Set to a percentage of current thresholds" on page 153).
- For key thresholds, a specific value that you specify, based on your knowledge of what is appropriate (using *Emergency Setup*; see "Set using Emergency Setup" on page 154).
- A value that you specify for a specific threshold.

These minimum thresholds are not necessarily hard limits. When traffic exceeds one of these thresholds, FortiDDoS uses its estimated threshold to detect and block rate anomalies until the observed traffic rate falls below the configured minimum threshold or rises to reach the maximum threshold.

Many of the traffic statistics graphs that you use the *Monitor* menu to access display the configured minimum threshold as a reference.

Estimated thresholds

Unlike most statically configured protection systems, FortiDDoS dynamically adjusts to the traffic patterns of the systems it protects. The FortiDDoS continuously and automatically calculates thresholds for multiple traffic parameters. It uses algorithms that can distinguish attack traffic from traffic increases that are the result of legitimate users accessing the protected system.

Although FortiDDoS generates these estimated thresholds periodically, it only uses them when they are higher than the configured minimum threshold. (The threshold limits that FortiDDoS

applies are never lower than the configured minimum thresholds that you configure using the methods described in “Configured minimum thresholds” on page 25). When this happens, the estimated threshold becomes a hard limit and FortiDDoS blocks any traffic that exceeds it. However, the estimated threshold changes periodically to account for any increase in legitimate traffic.

The minimum value of an estimated threshold is the configured minimum threshold. FortiDDoS calculates the maximum value of an estimated threshold using the specified adaptive limit values (see “Adaptive limit” on page 27).

Figure 5 shows traffic history in which FortiDDoS dynamically changes the threshold for TCP SYN Traffic in response to increases in traffic. The green line is the estimated threshold. Using past history, FortiDDoS adjusts the estimated threshold to a value that is often higher than the configured minimum threshold (the Threshold value shown in the top-left corner). The estimated threshold rises and drops dynamically as the traffic pattern (the blue line) changes. Traffic levels can exceed the configured minimum threshold without being blocked, but only within an allowed deviation.

Figure 5: Dynamic threshold changes

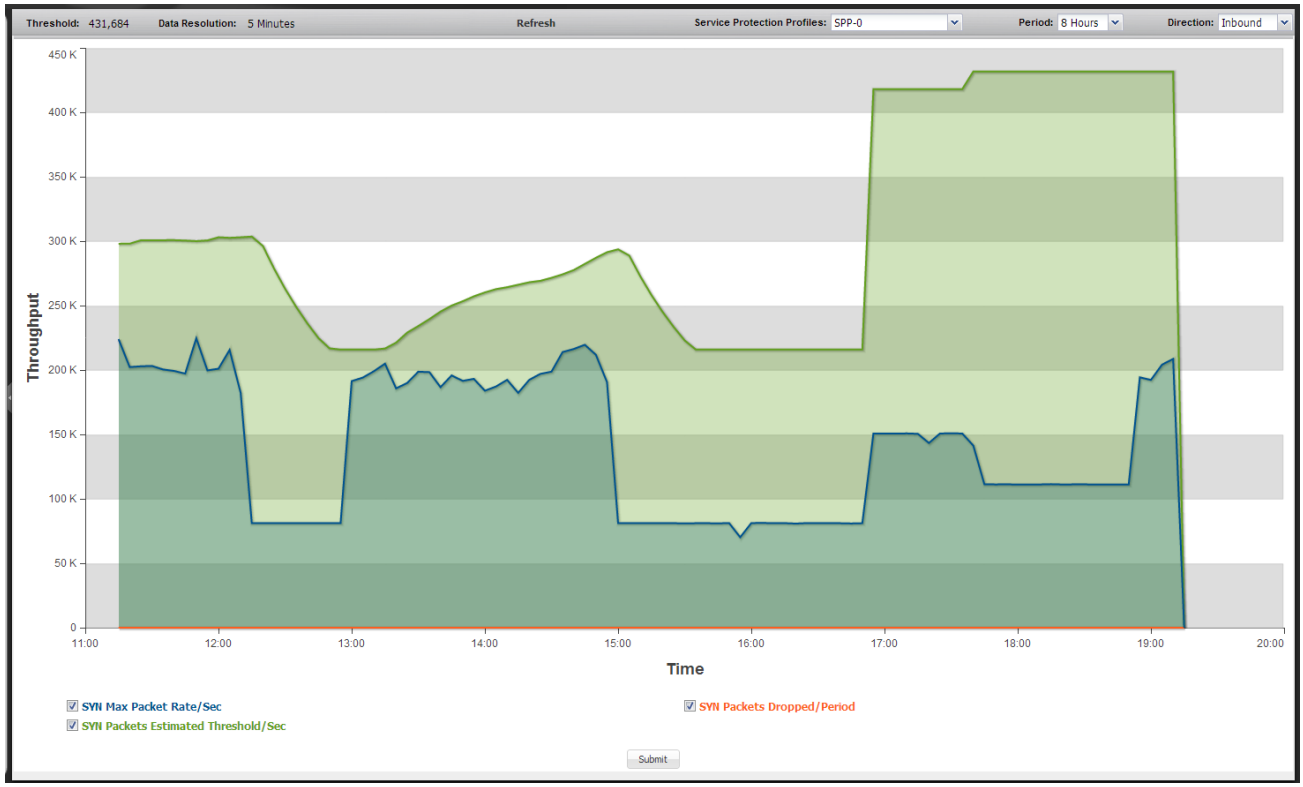

The following traffic statistics graphs display the estimated threshold for an individual traffic parameter:

- *Layer 3* graphs
  - *Most Active Source*
  - *Most Active Destination*
  - *Fragmented Packets*
- *Layer 4* graphs
  - *SYN Packets*
  - *SYN Per Source*
  - *SYN Per Destination*
  - *Connection Per Destination*
  - *ACK Per Destination*
  - *RST Per Destination*
  - *FIN Per Destination*
  - *ESTAB Per Destination*
  - *New Connections*
- *Layer 7* graphs
  - *HTTP Methods*
  - *Invites*
  - *Registers*
  - *Concurrent Invites*
  - *SIP Method*

Adaptive limit

When the estimated threshold is greater than the configured minimum threshold, FortiDDoS uses the estimated threshold value. However, FortiDDoS also calculates a maximum value for the estimated threshold. This upper limit is called the adaptive limit and is an upper rate limit beyond which all traffic is blocked.

FortiDDoS calculates the adaptive limit as a percentage of the configured minimum threshold. For example, if the adaptive limit is 150% (the default), FortiDDoS can use its dynamic threshold estimation algorithm to increase the calculated threshold to up to 150% of the value of the configured minimum threshold. An adaptive limit of 100% means no dynamic threshold estimation adjustment takes place once the configured minimum threshold is reached (that is, the threshold is a fixed value).

The adaptive limit percentage value used is for traffic across all layers.

Figure 6: Adaptive, minimum and fixed thresholds

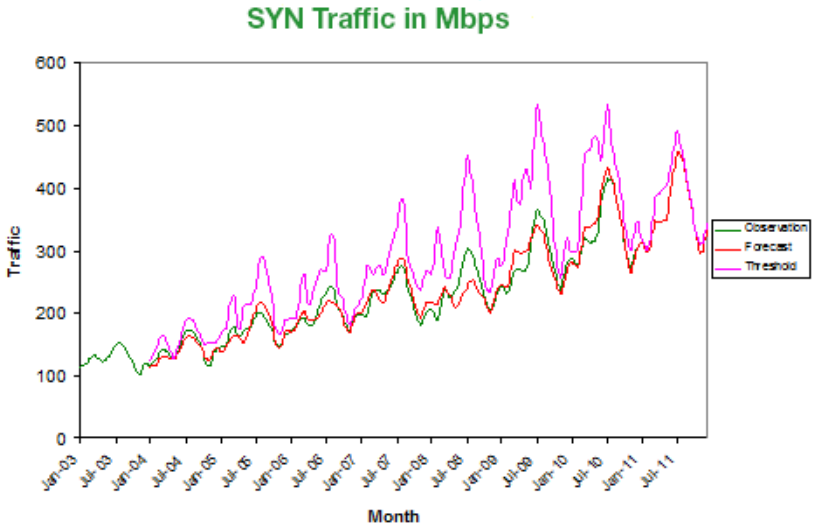

There are scenarios where FortiDDoS may drop legitimate traffic because it cannot adapt quickly enough to a sudden change in traffic patterns. For example, when a news flash or other important announcement increases traffic to a company's web site.

In these situations, you can increase all configured thresholds by a specific percentage. For more information, see "Set to a percentage of current thresholds" on page 153.

Hierarchical nature of protocols & implication on thresholds

FortiDDoS helps you determine the existing packet rates for different layers and their attributes, so that you don't have to guess these numbers. But if you are setting the thresholds manually, it is important to understand the protocol hierarchy that FortiDDoS uses to determine thresholds.

Figure 7: Protocol hierarchy for determining thresholds

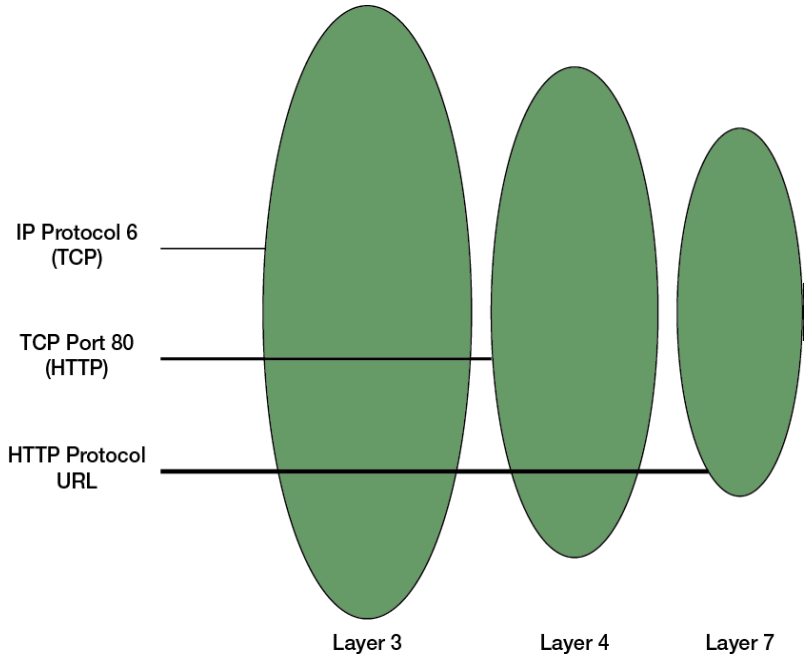

FortiDDoS analyzes thresholds at Layer 3, 4, and 7, any of which can trigger blocking. Remember that an HTTP packet is also a layer 4 (TCP) packet, an IP packet, and an Ethernet frame. It must be within the thresholds of all these meters in order to pass through the FortiDDoS gateway.

If you are setting thresholds manually, you need to be aware of the hierarchy of these meters so upper layer thresholds are not invalidated by lower layer limitations. (For example, so that HTTP packet rates are not unduly restricted by Layer 3 thresholds.)

The following image shows the hierarchy for determining thresholds on the FortiDDoS.

Figure 8: HTTP packet properties

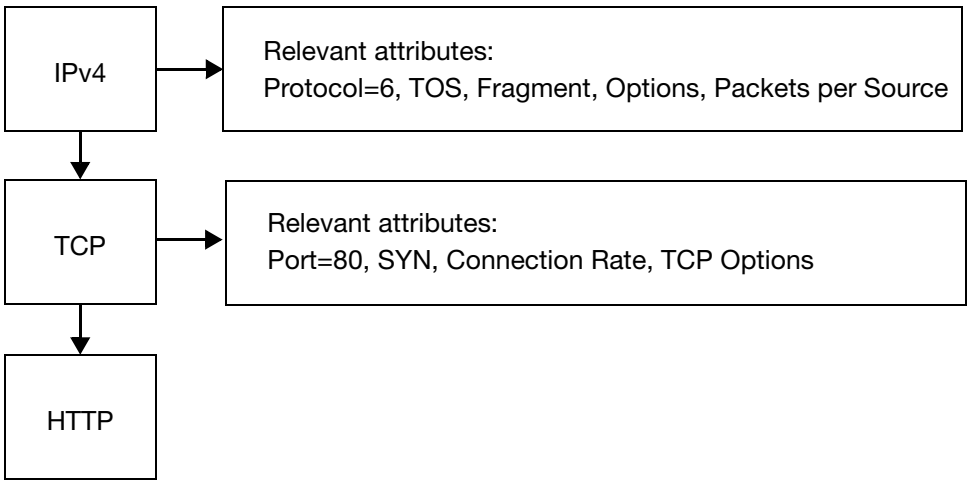

The above figure shows what happens when an HTTP packet passes through the FortiDDoS.

This packet is an IPV4 packet and TCP packet as well.

An IPv4 packet has following properties:

- Protocol
- TOS
- Fragment or not a fragment
- A Source IP address, and hence Packet rate from that specific source

A TCP packet has following properties:

- Destination Port
- SYN or not a SYN packet
- Connection Tuple. Since TCP packets belong to a connection, hence packet rate within that connection.

An HTTP packet has following properties:

- Method (for example, GET)
- URL

The following image shows the path of a UDP packet through the FortiDDoS.

Figure 9: UDP packet properties

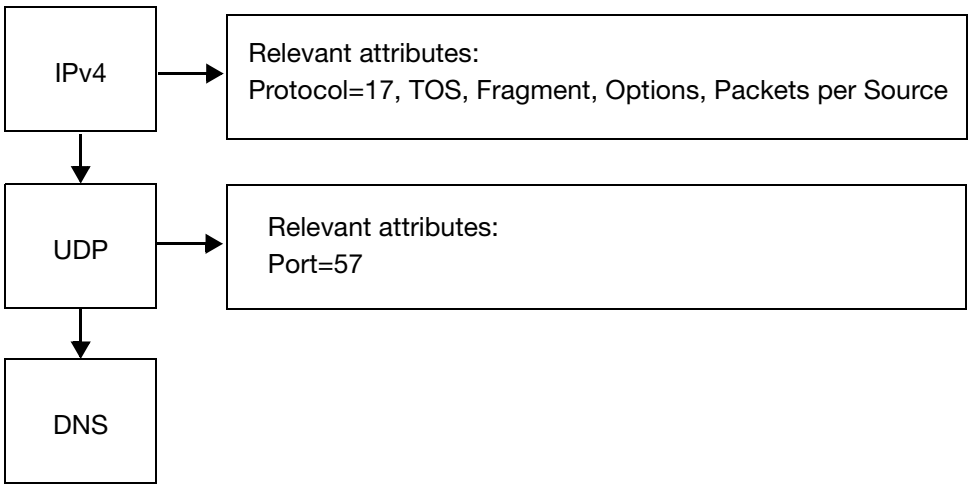

Figure above shows what happens when a UDP packet passes through the FortiDDoS.

This packet is an IPV4 packet and UDP packet as well.

An IPV4 packet has following properties:

- Protocol
- TOS
- Fragment or not a fragment
- IP Option values
- A source IP address, and hence Packet rate from that specific source

A UDP packet has following properties:

- Destination Port

If a server supports TCP, UDP, and ICMP services, the rate of IP packets is an aggregate of rates for TCP, UDP and ICMP packets. Similarly if the same server is a web server as well as an SMTP server, the TCP packet rate is the sum of port 80 and port 25 packet rates.

To summarize, because determining thresholds is a hierarchical process, avoid setting low thresholds on common conditions that can cause FortiDDoS to block legitimate traffic as well as attack traffic. The more specific you are about the type of traffic you want to allow as 'normal', the more effective the FortiDDoS is in blocking other traffic.

Granularity of traffic, corresponding rates and thresholds

FortiDDoS functions with set thresholds and limits. Some of the parameters used at various times are briefly described below. The benefits of the large number of thresholds settings that FortiDDoS uses - its granularity - are also discussed.

Keep in mind that packets and frames are measured against all applicable thresholds, not just the highest layer thresholds. For example, each HTTP (web request) packet is a TCP packet, but it is also an IP packet. The FortiDDoS blocks the packet if it violates any threshold at layer 3, 4 or 7. For more information, see "Hierarchical nature of protocols & implication on thresholds" on page 28.

Benefits of granularity

Denial of service attacks are mounted by professionals using botnets: networks of compromised machines. Increasingly, instead of simple bandwidth attacks, attackers try to avoid detection by creating attacks that mimic the behavior of a large number of clients. Because the content of the malicious requests does not differ from that of legitimate ones, the resulting attacks are hard to defend against using standard techniques.

Dodging an NBA system is very easy if attackers do coarse-grained rate-based control. In contrast, FortiDDoS has many thresholds at layer 3, 4, and 7. The granularity of thresholds makes it difficult for an attack to go undetected. With granular analysis, the appliance can drop only specific attack packets while the legitimate traffic can go through.

For example, if a few TCP connections are exceeding bandwidth, the system stops those connections from going forward rather than stop all connections. If a single destination is under attack, FortiDDoS stops packets to that destination while all others continue. During fragmented flood attacks, all non-fragmented packets continue as usual. During a port flood to a non-service port, the packets to other ports continue.

Granularity helps to increase the goodput (the throughput of useful data) of the system.

For a list of the traffic and counters at different layers, see Table 1 on page 20. For more information about individual traffic thresholds, see "Adjusting thresholds individually" on page 155.

FortiDDoS continuously learns rates in each direction for each of these traffic parameters. By setting thresholds granularity and adaptively in each direction for each of these parameters, FortiDDoS can easily thwart many of the rate-based attacks.

Using ACLs to block known attacks

FortiDDoS provides access control lists (ACLs) to prevent known attacks. For example, in a data center environment, the FortiDDoS ACLs can protect the router from getting overloaded by attacks from known flood traffic.

Because the ACLs are part of the core hardware architecture, they do not add to latency through the device when you enable or disable them.

The FortiDDoS ACLs provide a high level of granularity that is normally not found on traditional network equipment.

FortiDDoS has the following ACL features:

- A global ACL that applies to all traffic
- An ACL that denies access to specific subnets for specific protocols
- An ACL for each Service Protection Profile (SPP)
- An optional IP Reputation list provided through FortiGuard

The following table lists the traffic parameters you can add to an ACL.

Table 2: Available ACL parameters

Layer	Parameter	ACL
Layer 3	Any protocol (up to 256)	profile
	Fragment	profile
	IP netmask or address (up to 1 Million)	profile, global
	Geolocation (countries and regions), anonymous proxy, satellite provider	global
	IP-reputation (based on data from external public sources)	IP Reputation (subscription)
Layer 4	TCP port (up to 64k)	profile
	UDP port (up to 64k)	profile
	ICMP type/code (up to 64k)	profile
Layer 7	URLs (up to 32 767)	profile
	Host (512)	profile
	Referer (512)	profile
	Cookie (512)	profile
	User-Agent (512)	profile

See also

- Access control lists (ACLs)
- FortiGuard IP Reputation Service

Proxy IP addresses

It is common for an IP address to act as proxy—it is an intermediary for other IP addresses. Because the normal behavior of a proxy IP address can resemble an attack, you can configure FortiDDoS to automatically increase specific thresholds for sources that it determines are proxies.

For more information, see “Enabling higher thresholds for proxy server IP addresses” on page 132.

Analyzing, interpreting, & preventing attacks

This section explains DDoS attack events and discusses how to analyze and respond to attacks by interpreting the graphs.

A DDoS attack event is any condition that causes the FortiDDoS to block traffic when it is running in prevention mode.

Attack analysis workflow

Figure 10: Typical workflow for identifying and analyzing attacks

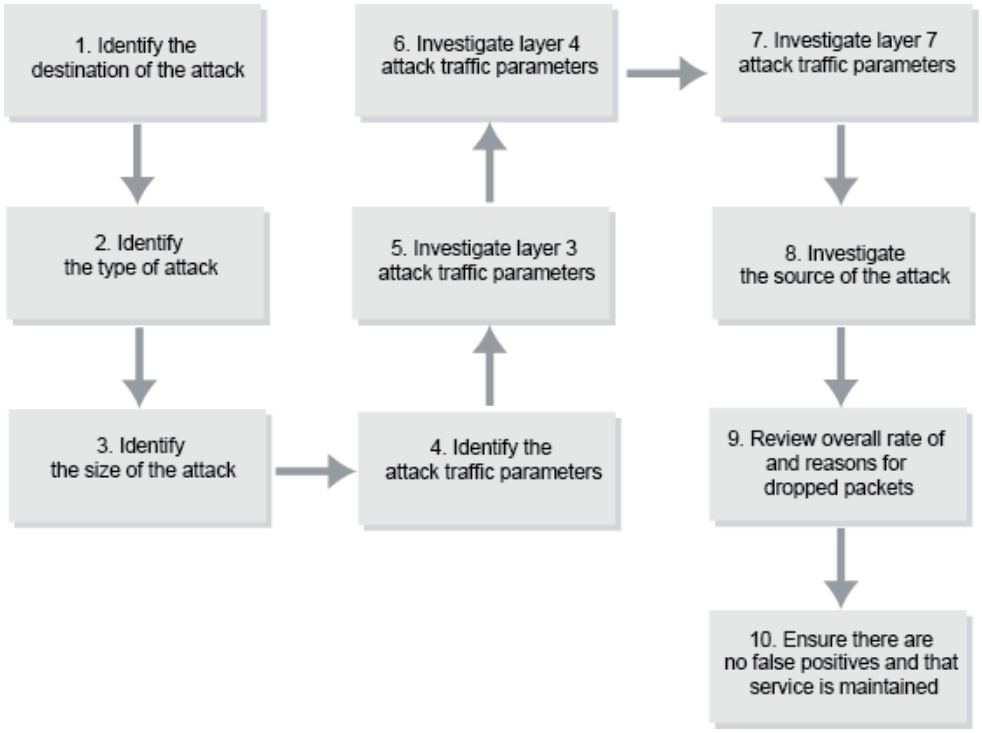

This flowchart is a general outline for investigating attacks. Typical investigations also include the following tasks:

- Checking the overall traffic level (see "Port Statistics graphs" on page 208)
- Checking aggregated flood drop graphs for all layers and individual layers (see "Working with graphs: Aggregate Flood Drops" on page 205)
- Identifying a distributed attack by checking the count of unique sources (see "Count of Unique Sources widget" on page 201 and "Layer 3 graphs" on page 219)
- Using Most Active Source, SYN per Source, and Connections per Source graphs to determine if a small number of sources is responsible for the anomalous traffic (see "Layer 3 graphs" on page 219 and "Layer 4 graphs" on page 220)
- Identify a distributed TCP attack by checking the count of TCP connections (See "Layer 4 graphs" on page 220)

The *Log & Report* menu provides access to overview information that includes:

- Which Service Protection Profile (SPP) is under attack
- The current top attack
- Which subnet is under attack
- The current top attackers

Complete information about the *Log & Report* menu items and the dashboards, graphs, and reports is provided in "Monitoring attack activity and other system information" on page 195).

Effects of crossing a threshold

FortiDDoS forwards packet traffic until it exceeds the threshold of a specific parameter. Traffic that exceeds the threshold is blocked for the configured blocking period. When blocking period is over, the threshold is checked again. The following examples assume that the blocking period has the default value of 15 seconds.

Example 1: A network receives too many packets with a specified protocol

- FortiDDoS drops incoming packets with the protocol that are destined for a specific network (specified as a subnet) for 15 seconds. It forwards all other packets.
- FortiDDoS tracks the source of the packets to determine if this is a single-source attack.
- After 15 seconds, FortiDDoS checks the rate of the packets against the threshold again.

Example 2: Too many mail messages coming to an SMTP server

- FortiDDoS drops incoming TCP packets destined for port 25 on the mail server (or the mail server's network) for 15 seconds. It forwards all other packets.
- FortiDDoS tracks the source of the packets to determine if this is a single-source attack. If there is a single source, the appliance blocks all packets from that source for 15 seconds.
- After 15 seconds, FortiDDoS checks the rate of the packets against the threshold again.
- Mail clients assume that the network is slowing down because TCP packets are lost. The clients start to send packets at a slower rate. No mail messages are lost.

Example 3: A web server receives too many SYN packets

- FortiDDoS checks TCP SYN packets destined for the web server. If they come from an IP address in the legitimate IP address table, FortiDDoS permits them to continue to the web server. The appliance allows these packets as long as their rate is lower than the new-connections threshold (designed to indicate zombie floods). FortiDDoS forwards all other SYN packets.
- If the IP address does not exist in the legitimate IP address table, and if the *SYN flood mitigation method* is *SYN Cookie*, FortiDDoS performs a proxy three-way handshake to validate the IP address.
- After 15 seconds, FortiDDoS checks the packet rate against the threshold again.

Example 4: A source has excessive concurrent connections

- If there are too many concurrent TCP connections from a single source, FortiDDoS blocks new connections until the number of concurrent connections is less than the threshold.
- Once the concurrent connection count goes down, FortiDDoS allows the source to establish new connections.
- FortiDDoS tracks the source of the connections to determine if this is a single-source attack. If there is a single source, the appliance blocks all packets for 15 seconds.
- After 15 seconds, FortiDDoS checks the connection rate against the threshold again.

Reducing false positives

When FortiDDoS blocks packet traffic because it exceeds the threshold of a specific traffic parameter, it blocks any traffic with the offending characteristic. As a result, during the blocking period, FortiDDoS may block traffic from legitimate sources in addition to traffic from a malicious source. FortiDDoS uses the following mechanisms to minimize the impact of these false positives:

- Because the blocking period is short (1 to 15 seconds), FortiDDoS frequently checks to see if the packet traffic no longer exceeds the threshold that detected the attack.
- FortiDDoS simultaneously attempts to determine whether the attack is not spoofed and can be attributed to one or a few sources. If it can identify these sources (called source attackers), it applies a penalty to them. The penalty makes these sources more likely to exceed the most active source threshold, which causes FortiDDoS to apply a longer blocking period.
- If it identifies attackers, FortiDDoS can stop blocking traffic from legitimate sources as soon as the standard, shorter blocking period is over, but continue to block traffic from source attackers for a longer period.

The following timeline illustrates how FortiDDoS can respond immediately to attacks but then adjust its attack mitigation activity to packets from specific sources only.

In this example, the standard blocking period is 15 seconds and the blocking period for source attackers is 60 seconds (the default value):

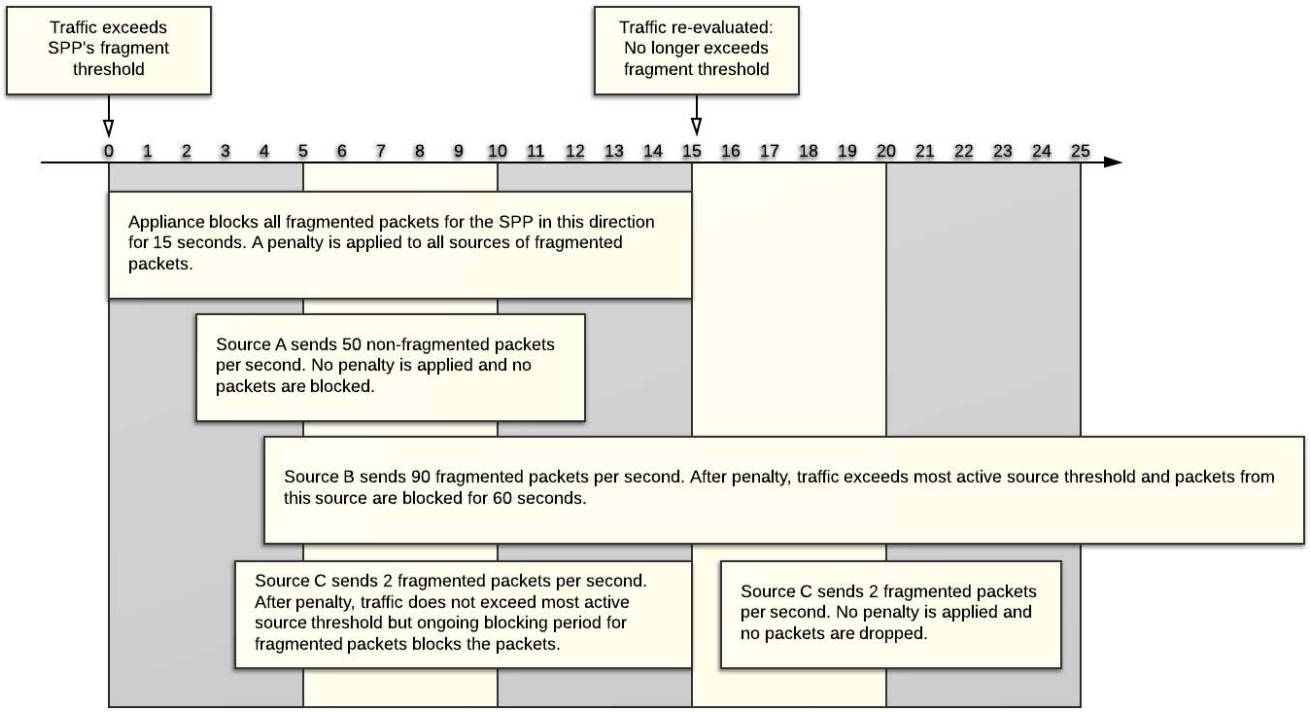

The penalty factor for source attackers is 16, and the most active source threshold is 100 packets per second. So when Source B sends 90 fragmented packets, the calculated rate is 1440 packets per second, which exceeds the most active source threshold. But when Source C sends 2 fragmented packets per second, the calculated rate of 32 packets per second does not exceed the threshold. Thus, FortiDDoS applies the longer blocking period to Source B only.

Source C, which sends an insignificant number of fragmented of packets, is blocked only for the length of the shorter, standard blocking period.

Events by layer

For descriptions of events that FortiDDoS records in the DDoS Attack Log, see “DDoS Attack Log and DDoS Subnet Attack Log” on page 222.

Analyzing attacks

The first indication that FortiDDoS has detected an attack is the dashboard. The dashboard widgets summarize the number of packets dropped and a count of unique sources.

If the attack lasts for 5 minutes or more, you can see a graphical representation of it in one or more places. In many cases, the first place you look is *Monitor* > *Aggregate Flood Drops*, which displays which packets have been dropped at each layer. This graph allows you to further refine

your search to Layer 3, Layer 4 or Layer 7 parameters. The table below shows some well-known attack scenarios and how they may be manifested on the FortiDDoS

Table 3: Causes, Diagnosis and Prevention of Attacks

Attack	Description	Threshold to monitor/adjust	Events to watch
SYN attack	An excessive number of packets on a specific TCP Port. In most cases, the source address is spoofed.	Layer 3 - TCP protocol (6) Layer 4 - TCP ports on which the server is listening and ports that are allowed by the firewall and ACL Layer 4 - SYN	Protocol 6 Flood SYN Flood Zombie Flood Port Flood
UDP flood attack	An excessive number of UDP packets.	Layer 3 – UDP protocol (17) Layer 4 – UDP ports on which the server is listening and ports which are allowed by the firewall and ACL	Protocol 17 Flood Port Flood
ICMP flood	An excessive number of ICMP packets.	Layer 3 – ICMP protocol (1) Layer 4 – ICMP type and codes combinations that are allowed by the firewall and ACL.	Protocol 1 Flood Layer 4 ICMP Flood of a specific type and code
Fragment flood	An excessive number of fragmented packets.	Layer 3 – Fragmented packets	Fragment Flood
Source flood	A single source sends excessive number of IP packets.	Layer 3 – Most active source	Source Flood
Zombie attack	Too many legitimate IP sources send legitimate TCP packets.	Layer 3 – TCP protocol (6) Layer 4 – TCP ports on which the server is listening and ports that are allowed by the firewall and ACL. Layer 4 – SYN Layer 4 – Established connections per destination (estab-per-dst) Layer 4 - SYN per source (syn-per-src)	Layer 3 Protocol 6 SYN Flood Zombie Flood Port Flood SYN Flood from Source

Table 3: Causes, Diagnosis and Prevention of Attacks

Attack	Description	Threshold to monitor/adjust	Events to watch
Slow connection buildup	Legitimate IP sources send legitimate TCP connections but do it slowly and remain idle, which fills up the server's connection table memory.	Layer 3 – TCP protocol (6) Layer 4 – TCP ports on which the server is listening and ports that are allowed by the firewall and ACL. Layer 4 – SYN Layer 4 – New connections Layer 4 - Concurrent connections per source Layer 4 - Concurrent connections per destination	Layer 3 Protocol 6 SYN Flood Zombie Flood Port Flood Excess Concurrent Connections/ Source Excess Concurrent Connections/ Destination
Slammer attack	An excessive number of packets on UDP Port 1434.	Layer 3 – UDP protocol (17) Layer 4 – UDP port 1434	Protocol 17 UDP Flood Port Flood –1434
DNS attack	An excessive number of packets on UDP Port 53.	Layer 3 - UDP protocol (17) Layer 4- UDP port 53	Protocol 17 UDP Flood UDP Port 53 Flood ICMP Port/Host not available Flood
MyDoom attack	Excessive number of packets on HTTP from zombies.	Layer 3 – TCP protocol (6) Layer 4 – TCP port 80 Layer 4 – SYN Layer 4 – New connections Layer 4 – Established Connections	Protocol 6 Flood SYN Flood Zombie Flood Port Flood
Smurf attack	Traffic that appears to originate from the target server's own IP address or somewhere on its network. Targeted correctly, it can flood the network with pings and multiple responses.	Layer 3 – ICMP protocol (1) Layer 4 – ICMP type and codes combinations that are allowed by the firewall and ACL.	Protocol 1 Flood ICMP Flood of Echo-Request/ Response Type (Type= 0, Code = 0)

Table 3: Causes, Diagnosis and Prevention of Attacks

Attack	Description	Threshold to monitor/adjust	Events to watch
Fraggle attack	Spoofed UDP packets to a list of broadcast addresses. Usually the packets are directed to port 7 on the target machines, which is the echo port. Other times, it is directed to the Character Generator Protocol (CHARGEN) port. Sometimes a hacker is able to set up a loop between the echo and chargen port. FortiDDoS has applicable thresholds for each Service Protection Profile (SPP) at layer 3 and 4. A layer 4 packet is, by definition, also a layer 3 packet. However, blocked traffic is only displayed at the highest layer at which a threshold was violated. For example, FortiDDoS identifies a packet causing a TCP connection flood to be layer 4 dropped traffic only, even though the corresponding layer 3 envelope is also blocked.	Layer 3 – ICMP protocol (1) Layer 3 – UDP protocol (17) Layer 4 – UDP echo port (7) Layer 4 – Daytime Protocol port (13) Layer 4 – Quote of the Day (QOTD) port (17) Layer 4 – UDP Character Generator protocol (CHARGEN) (19) Layer 4 – ICMP Type/Codes specific to host/port not available.	Protocol 1 Flood Protocol 17 Flood UDP Port 7 Flood UDP Port 13 Flood UDP Port 17 Flood UDP Port 19 Flood ICMP Flood of Port Not Available Type, Code (3,3) ICMP Flood of Host Not Available Type, Code (3,1)
HTTP GET attack	Excessive number of packets on HTTP from zombies.	Layer 3 – TCP protocol (6) Layer 4 – TCP ports on which the server is listening and ports that are allowed by the firewall and ACL. Layer 4 – SYN Layer 4 – New Connections Layer 4 - Concurrent connections per source Layer 4 - Concurrent connections per destination Layer 7- HTTP Methods Layer 7 - URL	Protocol 6 Flood SYN Flood Zombie Flood Port Flood TCP Connection Flood HTTP OpCode Flood (HTTP Method Flood) URL Flood

See also

- Adjusting thresholds
- Traffic graphs
- Logging
- Reports

Working with attack reporting

The FortiDDoS DDoS attack reporting features are designed to maximize the processing resources that are available for preventing attacks. By restricting reporting to key traffic parameters, FortiDDoS is better able to withstand multi-gigabyte attacks. However, as a result, reporting tools such as the DDoS attack log do not always include detailed traffic parameter information. Outside of specific scenarios, FortiDDoS does not report source and destination IPs and ports, protocols, and so on, for every dropped or blocked packet.

For example, it is not uncommon for a FortiDDoS monitoring and regulating a 1 Gbps traffic flow to be the target of a 700 Mbps SYN flood for 8 hours. If FortiDDoS stored every source and destination IP, port, and protocol, the logging demands (via hard disk, syslog, or SNMP trap) would soon overwhelm the disk or network.

By concentrating its resources on dropping attack traffic and maintaining service, FortiDDoS allows you to focus your attention elsewhere and still provide you with helpful and relevant information when an attack is underway.

The rest of this section discusses some of the questions that users often have about the attack log events.

Why does FortiDDoS not report the destination for some types of attacks?

To keep its reporting processes manageable, FortiDDoS does not always report a source or destination IP in the DDoS attack log. For example, for a HTTP GET flood, FortiDDoS reports the protocol of the dropped packets but not their destination IP or port.

To determine the destination of an attack, use the service protection profile (SPP) policy configuration to organize your network into subnets that you specify. Then use reporting tools such as *Log & Report* > *Report Browse* to see which subnet is under attack.

For more information, see “Identifying IP addresses and subnets to protect (SPP creation)” on page 111 and “DDoS Attack Activity report types” on page 255.

Why does FortiDDoS not report the source of a SYN flood?

During a SYN flood attack, FortiDDoS reports a SYN flood, identifies the service protection profile (SPP) that is under attack, and reports how many packets it has dropped. SYN floods are spoofed — the reported source of the packets is not their true source. Trying to determine the source IP or report potentially millions of source IPs that have no consistent pattern is resource-intensive and does not help you to determine the identity of the attacker.

Why does FortiDDoS report some attack events every 5 minutes instead of when they happen? Why does it not generate an audible alert or other type of alarm?

For some types of attacks, such as TCP and ICMP checksum errors, FortiDDoS collects aggregate data and reports every 5 minutes only. If the appliance reported each dropped packet

as soon as it dropped it and generated some kind of alert every time it dropped a packet, it would log events and generate alerts continuously.

Service Protection Profiles (SPPs)

Virtualization is at the core of the FortiDDoS product. The virtualization feature brings down the cost of protection per entity by providing a replicated set of counters and thresholds. It enables a single FortiDDoS to protect multiple network zones with the same granularity.

You can use FortiDDoS to configure up to 7 protection zones in addition to the default zone. Set up zones for specific departments, geographic locations or functions within an organization. You can then set security policies for each Service Protection Profile. Virtualization allows you to separate a single physical FortiDDoS device into up to 8 logical devices based on IP address and mask, right down to an individual host, if required.

Figure 11: Network with FortiDDoS protecting multiple subnets

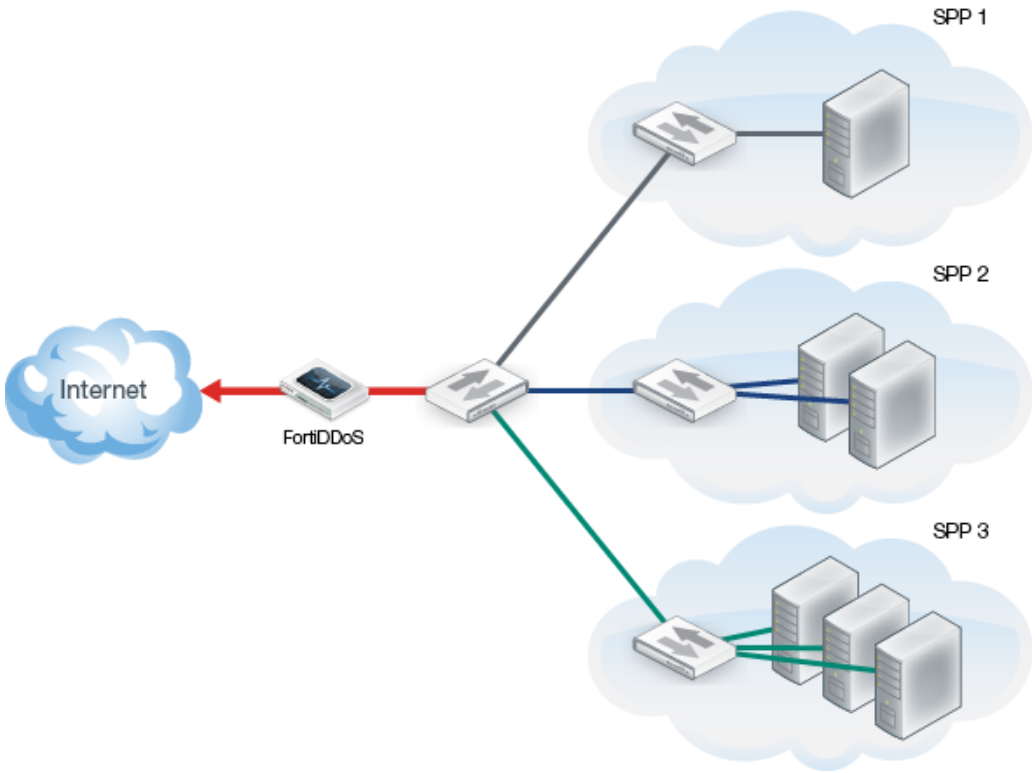

Benefits of virtualization

The configuration of each Service Protection Profile can be under the control of a different administrator. Profiles allow FortiDDoS to protect the system or systems using a customized configuration and traffic history, including:

- **ACLs —** The FortiDDoS ACL functionality allows a profile's administrator to permit or deny traffic based on conditions that are specific to the profile. For example, if a profile contains

subnets that contain only file and print servers, then the administrator of the profile could deny all HTTP traffic.

- **Thresholds —** Profile administrators can set thresholds that are based on conditions that are specific to the profile. For example, threshold for a server with a normal traffic rate of 1 Mbps for HTTP has different thresholds than a server with a rate of 100 Mbps.
- **Event Monitoring, Notifications and Reporting —** FortiDDoS can identify, prevent, monitor, notify and report attacks on individual profiles.
- **Support for managed service environments —** Because the Service Protection Profile functionality allows you to apply unique and independent security policies, it is perfect for the service provider environment where clients require customized services.

The Service Protection Profile functionality is implemented by IP address. By protecting a block of IP addresses utilizing Classless Inter-domain Routing (CIDR) blocks. CIDR-based Virtual IPs implementation allows granularity down to an individual host level.

See also

- Identifying IP addresses and subnets to protect (SPP creation)
- Service Protection Profile settings

How to use the web UI

This topic describes aspects that are general to the use of the web UI, a graphical user interface (GUI) that provides access the FortiDDoS appliance from within a web browser.

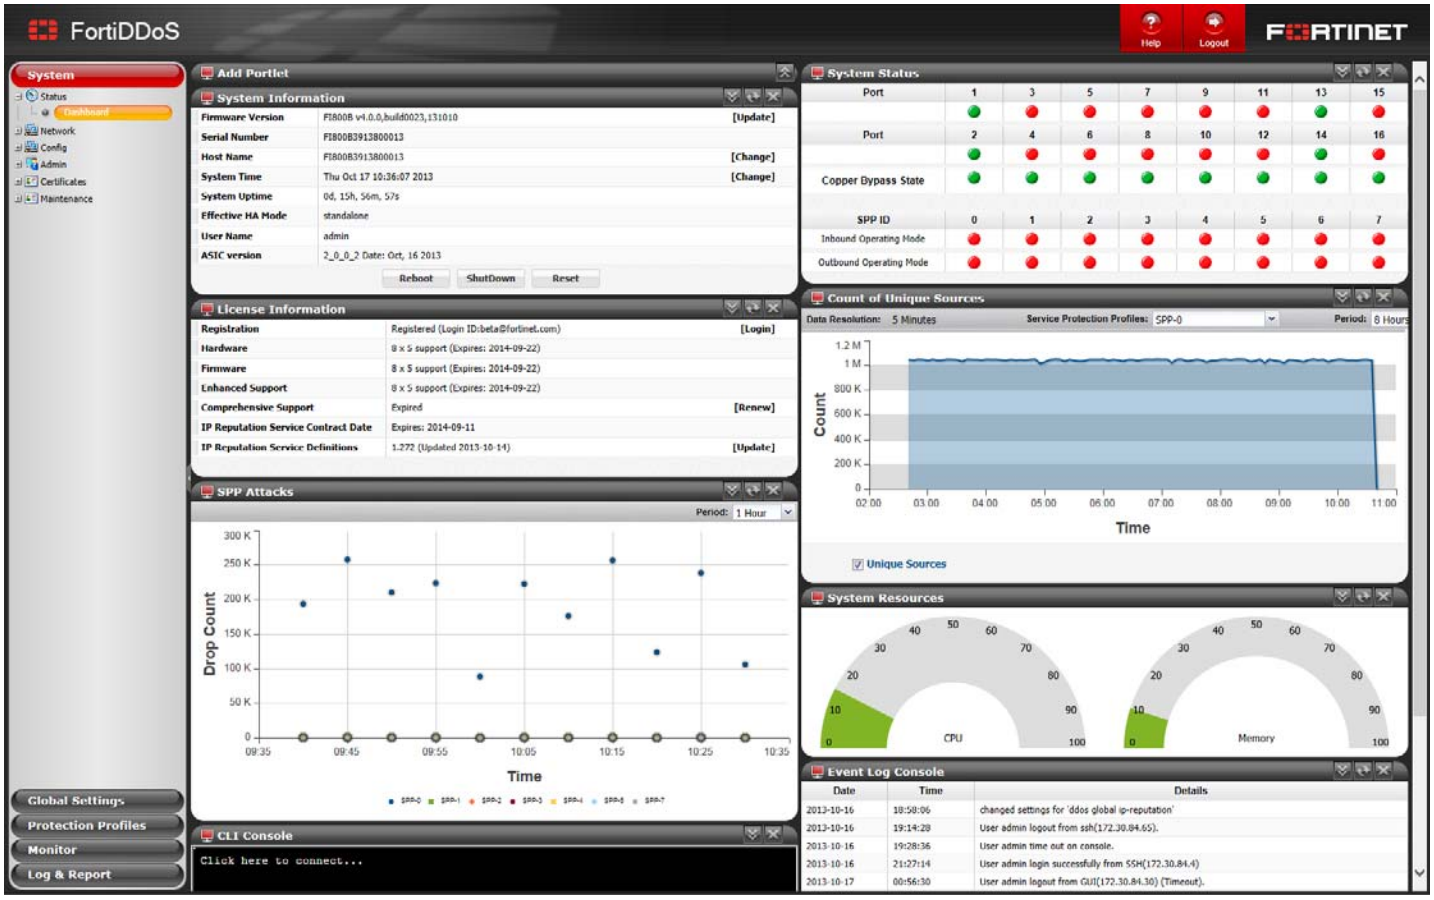

System requirements

The management computer that you use to access the web UI must have a compatible web browser, such as:

- Mozilla Firefox 20 or greater, or
- Microsoft Internet Explorer 8.0 or greater

To minimize scrolling, the computer's screen should have a resolution that is a minimum of 1280 x 1024 pixels.

URL for access

You access the web UI by URL, using a network interface on the FortiDDoS appliance that you have configured for administrative access.

For first-time connection, see "Connecting to the web UI" on page 76.

The default URL to access the web UI through the network interface on mgmt1 is:

https://192.168.1.99/

If the network interfaces were configured during installation of the FortiDDoS appliance (see “Configuring network interfaces, gateway, and DNS” on page 98), the URL and/or permitted administrative access protocols may no longer be in their default state. In that case, use either a DNS-resolvable domain name for the FortiDDoS appliance as the URL, or the IP address that was assigned to the network interface during the installation process.

For example, you might have configured port2 with the IP address 10.0.0.1 and enabled HTTPS. You might have also configured a private DNS server on your network to resolve fortiddos.example.com to 10.0.0.1. In this case, to access the web UI through port2, you could enter either `https://fortiddos.example.com/` or `https://10.0.0.1/`.

For information on enabling administrative access protocols and configuring IP addresses for the FortiDDoS appliance, see “Configuring network interfaces, gateway, and DNS” on page 98.

If the URL is correct and you still cannot access the web UI, you may also need to configure FortiDDoS to accept login attempts for your administrator account from that computer (that is, trusted hosts), and/or static routes. For details, see “Administrators” on page 171 and “Adding a gateway” on page 103.

Permissions

Depending on the account that you use to log in to the FortiDDoS appliance, you may not have complete access to all CLI commands or areas of the web UI.

Access profiles control which commands and areas an administrator account can access. Access profiles assign one of the following types of access to each area of FortiDDoS:

- *Read* (view access)
- *Read-Write* (view, change, and execute access)
- no access

For more information on configuring the access profile for an administrator account can use, see “Restricting permissions” on page 174.

Table 4: Areas of control in access profiles

Access profile setting	Grants access to*	
System	*System* > ...	Web UI

`system`	`config spp ...`
	`config system ...`
	`show spp ...`
	`show system ...`
	`show full-configuration`
	`diagnose debug ...`
	`diagnose hardware ...`
	`diagnose netlink ...`
	`diagnose sniffer ...`
	`diagnose sys top ...`
	`execute ...`
		CLI

Table 4: Areas of control in access profiles

Access profile setting	Grants access to*	
Global Settings	*Global Settings* >...	Web UI

`global-settings`	`config ddos global ...`
	`get system status`
	`get system performance`
	`show system status`
	`show system performance`
	`diagnose ...`
		CLI

Protection Profiles	*Protection Profiles* > ...	Web UI

`protection-profiles`	`get system status`
	`get system performance`
	`show system status`
	`show system performance`
	`show full-configuration`
	`diagnose ...`
		CLI

Monitor	*Monitor* > ...	Web UI

`monitor`	`get system status`
	`get system performance`
	`show system status`
	`show system performance`
	`show full-configuration`
	`diagnose ...`
		CLI

Log & Report	*System* > *Status* > *Dashboard* ... (statistics widgets only) *Log & Report* > ...	Web UI
`log`	`config system mailserver` `config log ...`	CLI

***** For each `config` command, there is an equivalent `get`/`show` command, unless otherwise noted.

`config` access requires write permission.

`get`/`show` access requires read permission.

Unlike other administrator accounts, the administrator account named admin exists by default and cannot be deleted. The `admin` administrator account is similar to a root administrator account. This administrator account always has full permission to view and change all FortiDDoS configuration options, including viewing and changing ***all*** other administrator accounts. Its name and permissions cannot be changed. It is the only administrator account

that can reset another administrator's password without being required to enter that administrator's existing password.

Set a strong password for the `admin` administrator account, and change the password regularly. By default, this administrator account has no password. Failure to maintain the password of the `admin` administrator account could compromise the security of your FortiDDoS appliance.

For complete access to ***all*** commands and abilities, you must log in with the administrator account named `admin`.

See also

- Restricting permissions
- Administrators
- Trusted hosts

Trusted hosts

As their name implies, trusted hosts are assumed to be (to a reasonable degree) safe sources of administrative login attempts.

Configuring the trusted hosts of your administrator accounts (*Trusted Host*) hardens the security of your FortiDDoS appliance by further restricting administrative access. In addition to knowing the password, an administrator must connect only from the computer or subnets you specify. The FortiDDoS appliance does not allow logins for that account from any other IP addresses. If all administrator accounts are configured with specific trusted hosts, FortiDDoS ignores login attempts from all other computers. This eliminates the risk that FortiDDoS could be compromised by a brute force login attack from an untrusted source.

Trusted host definitions apply both to the web UI and to the CLI when accessed through Telnet, SSH, or the *CLI Console widget*. Local console access is ***not*** affected by trusted hosts, as the local console is by definition not remote, and does not occur through the network.

See also

- Administrators
- Restricting permissions
- Permissions

Global web UI & CLI settings

Some settings for connections to the web UI and CLI apply to all administrator accounts.

To configure administrator settings

1. Go to *System* > *Admin* > *Settings*.

   To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *System* category. For details, see "Permissions" on page 44.

2. Configure these settings:

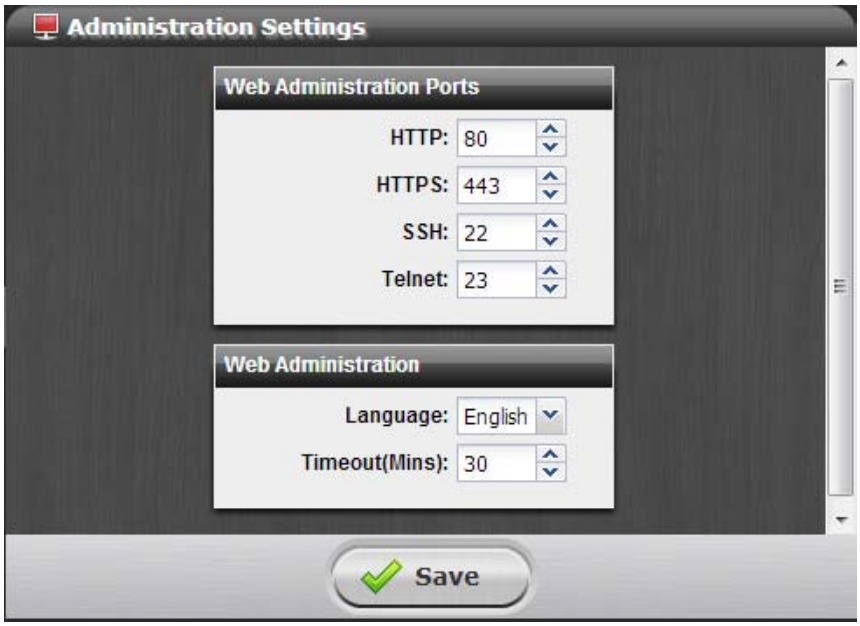

Setting	Description
Administration Ports	
HTTP	Type the port number on which the FortiDDoS appliance listens for HTTP administrative access. The default is 80. This setting has an effect only if *HTTP* is enabled as an administrative access protocol on at least one network interface. For details, see "Configuring the network interfaces" on page 98.
HTTPS	Type the port number on which the FortiDDoS appliance listens for HTTPS administrative access. The default is 443. This setting has an effect only if *HTTPS* is enabled as an administrative access protocol on at least one network interface. For details, see "Configuring the network interfaces" on page 98.
SSH	Type the port number on which the FortiDDoS appliance listens for SSH administrative access. The default is 22. This setting has an effect only if *SSH* is enabled as an administrative access protocol on at least one network interface. For details, see "Configuring the network interfaces" on page 98.
Telnet	Type the port number on which the FortiDDoS appliance listens for Telnet administrative access. The default is 23. This setting has an effect only if *TELNET* is enabled as an administrative access protocol on at least one network interface. For details, see "Configuring the network interfaces" on page 98.

Setting	Description
Web Administration	
Language	Select which language to use when displaying the web UI. Languages currently supported by the web UI are: • English • Simplified Chinese The display's web pages use UTF-8 encoding, regardless of which language you choose. UTF-8 supports multiple languages, and allows them to display correctly, even when multiple languages are used on the same web page. For example, your organization has web sites in both English and Simplified Chinese. Your FortiDDoS administrators prefer to work in the English version of the web UI. They can use the web UI in English while writing rules to match content in both English and Simplified Chinese ***without*** changing this setting. Both the rules and the web UI display correctly, as long as all rules were input using UTF-8. Usually, your text input method or your management computer's operating system should match the display by also using UTF-8. If they do not, your input and the web UI may not display correctly at the same time. For example, your web browser's or operating system's default encoding for Simplified Chinese input may be GB2312. However, you ***usually*** should switch it to be UTF-8 when using the web UI, ***unless*** you are writing regular expressions that must match HTTP client's requests, and those requests use GB2312 encoding. **Note:** This setting does ***not*** affect the display of the CLI.
Timeout	Type the number of minutes that a web UI connection is idle before FortiDDoS automatically logs out the administrator. The maximum value is 480 minutes (8 hours) and the minimum 1 minute. To maintain security, reduce the idle timeout from the default value of 30 minutes.

3. Click *Save.*

See also

- Configuring the network interfaces

Figure 12: Web UI parts

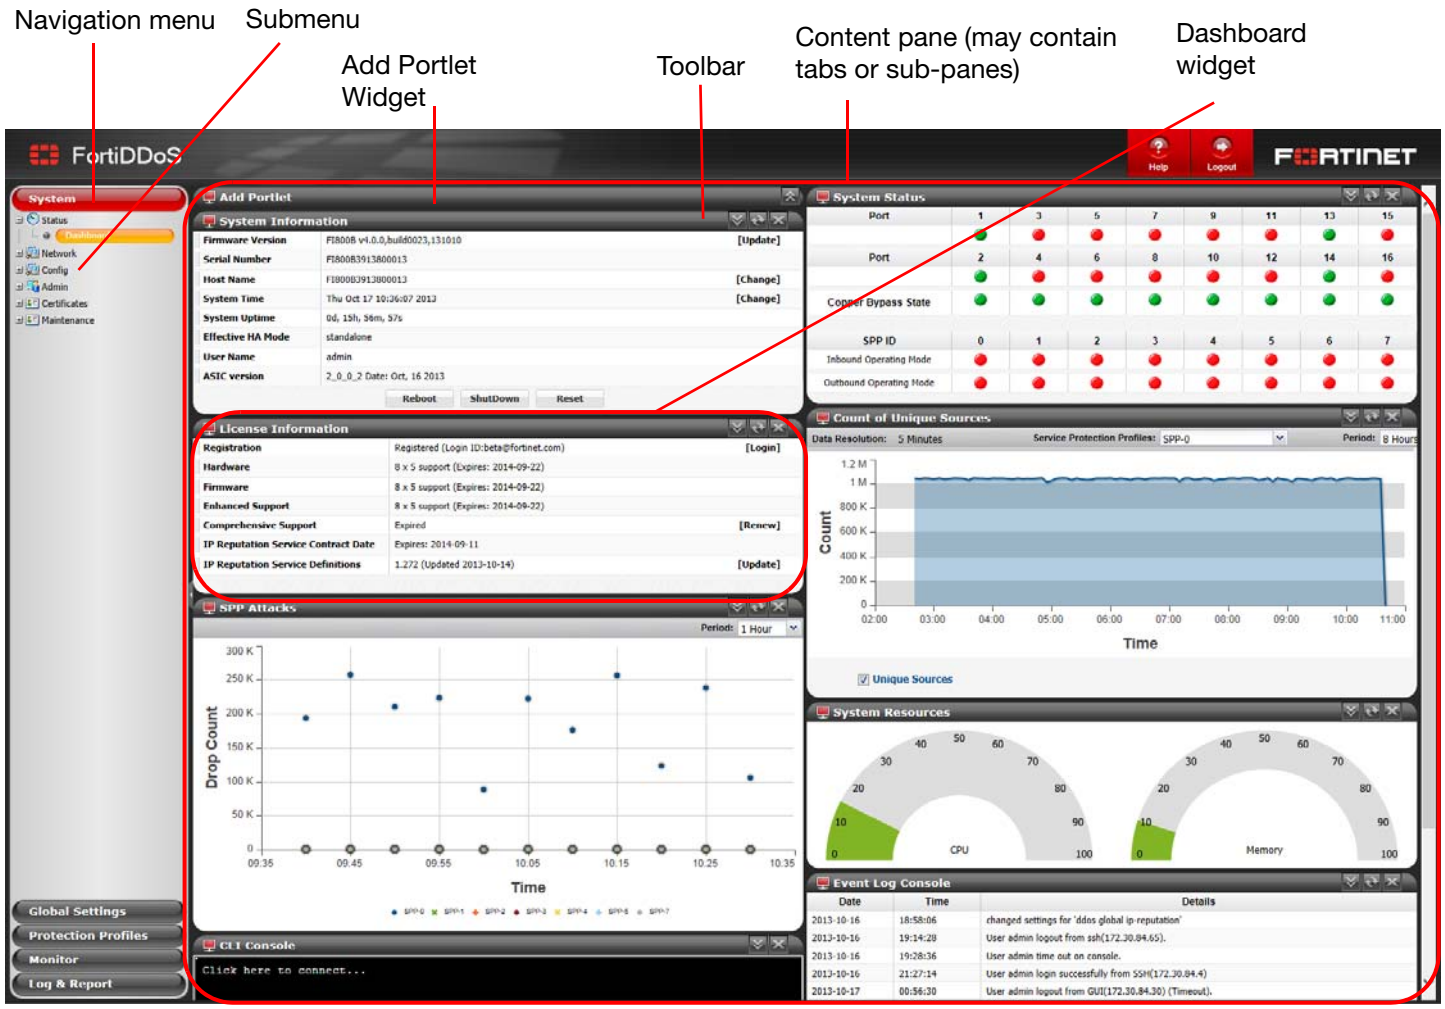

A navigation menu is located on the left side of the web UI. To expand a menu item, simply click it. To expand a submenu item click the + button located next to the submenu name, or click the submenu name itself. To view the pages located within a submenu, click the name of the page.

Do not use your browser's *Back* button to navigate — pages may not operate correctly. Instead, use the navigation menu, tabs, and buttons within the pages of the web UI.

To expand or collapse an area of the menu, click the name of the area itself. Within each area may be multiple submenus. To expand or collapse a submenu, click the + or - button next to the submenu name, or click the name of the submenu itself.

Within each submenu may be one or more tabs or sub-panes, which are displayed to the right of the navigation menu, in the content pane. At the top of the content pane is a toolbar. The toolbar contains buttons that enable you to perform operations on items displayed in the content pane, such as importing or deleting entries.

See also

- Deleting entries
- Renaming entries

Deleting entries

To delete a part of the configuration, you must first remove all references to it.

For example, an item in the list of global addresses (*Global Settings* > *Address* > *Access Config*) named “bogon_address” is the *Source address* value for an item in the global ACL (*Global Settings* > *Access Control List* > *Access Control List*) named “bogon_1”. You cannot delete “bogon_address” until you have deleted “bogon_1” and any other items that reference the address item.

Back up the configuration before deleting any part of the configuration. Deleted items cannot be recovered unless you upload a backup copy of the previous configuration. See “Backups” on page 167 and “Restoring a previous configuration” on page 169.

If you do not know where your configuration refers to the entry that you want to delete, to find the references, you can download a backup of the configuration and use a plain text editor to search for the entry’s name.

Predefined entries included with the firmware cannot be deleted.

See also

- Buttons, menus, & the displays
- Renaming entries

Renaming entries

In the web UI, each entry’s name is not editable after you create and save it.

For example, let’s say you create a SPP policy with the *Name* value “web_content”. While configuring the profile, you change your mind about the name a few times, and ultimately you change *Name* to “web_servers”. Finally, you click *Save* to save the policy. Afterwards, if you edit the policy, you can edit most of the settings. However, *Name* is unavailable and you cannot edit it.

While you cannot edit *Name*, you can achieve the same effect by other means.

To rename an entry

Alternatively, if you need to rename an item that is ***only*** referenced in the core configuration file, you can download a backup copy, use a plain text editor to find and replace the entry's old name, and then restore the modified configuration backup file to the appliance. Where there are many references, this may save time.

1. Clone the entry, supplying the new name.
2. In ***all*** areas of the configuration that refer to the old name, replace the old entry name by selecting the new name.

If you do not know where your configuration refers to the entry that you want to delete, to find the references, you can download a backup of the configuration and use a plain text editor to search for the entry's name.

3. Delete the item with the old name.

See also

- Buttons, menus, & the displays
- Deleting entries

Shutdown

Always properly shut down the FortiDDoS appliance's operating system ***before*** turning off the power switch or unplugging it. This causes it to finish writing any buffered data, and to correctly spin down and park the hard disks.

Do not unplug or switch off the FortiDDoS appliance without first halting the operating system. Failure to do so could cause data loss and hardware problems.

To power off the FortiDDoS appliance

1. Access the CLI or web UI. For details, see "Connecting to the web UI or CLI" on page 74.
2. From the CLI console, enter the following command:

   ```
   execute shutdown
   ```

   Alternatively, if you are connected to the web UI, go to *System* > *Status* > *Dashboard*, and in the *System Information* widget, click *Shut Down*.

   You may be able to hear the appliance become quieter when the appliance halts its hardware and operating system, indicating that power can be safely disconnected.

Figure 13: Turning off the system

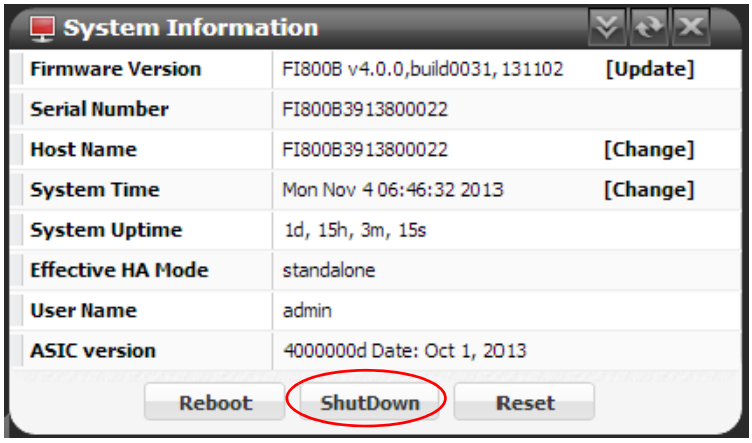

3. Disconnect the power cable from the power supply.

How to set up your FortiDDoS

These instructions describe the tasks you perform to initially add a FortiDDoS appliance to your network. These instructions assume that you have already read “Key concepts” on page 16 and are familiar with the fundamental concepts related to FortiDDoS devices.

For an overview of the initial configuration procedures, see “Configuration workflow” on page 21.

Registering your FortiDDoS

Before you begin, take a moment to register your Fortinet product at the Fortinet Technical Support web site:

https://support.fortinet.com

Many Fortinet customer services such as firmware updates, technical support, and FortiGuard services require product registration.

For more information, see the Fortinet Knowledge Base article Registration Frequently Asked Questions.

Planning the network topology

The FortiDDoS appliance is designed to protect a system or a network of systems from rate-based and anomaly attacks. If multiple systems or workgroups are protected by a FortiDDoS appliance, a switch is required between the FortiDDoS appliance and the protected systems.

See also

- Basic topology for FortiDDoS
- Basic web hosting deployment
- External bypass switches for maintenance & failover
- Load balancing
- Traffic diversion
- Topology for synchronizing the configuration of two FortiDDoS appliances
- Feature availability in a one-way traffic configuration

Basic topology for FortiDDoS

The FortiDDoS appliance is positioned 'inline', meaning it is installed between the one or more protected systems and the rest of the network. In the simple network shown in Figure 14, data passes through the FortiDDoS appliance as it travels to and from a protected system and the rest of an Ethernet local area network.

Because the FortiDDoS appliance is stateful and bidirectional, the data packet traffic is described as either incoming (inbound) and outgoing (outbound).

Figure 14: Logical network configuration in serial mode

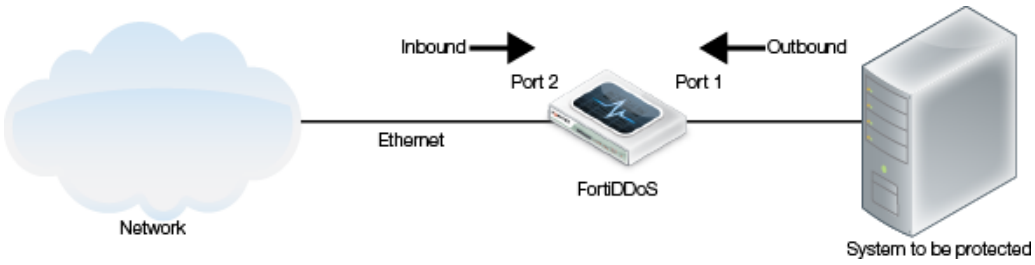

In this example, the Ethernet segment connected to the protected system is connected to Port 1. The Ethernet segment connected to the rest of the network (typically the Internet) is connected to Port 2.

For networks with multiple servers, you can simply dedicate a port pair to each server. For example, to protect 8 servers with connections with a total throughput of up to 8 Gbps, connect each server to a port pair on a FortiDDoS 800B. In more complex situations, you can add switches to distribute the load (see "Load balancing" on page 58).

You can add FortiDDoS appliances to the network in series (in-line) or place them in tandem with a bypass switch to avoid data path failures.

Ports 17-20 on the FortiDDoS 2000B connect to the network using fiber-optic cable via small form-factor pluggable (SFP) transceivers and have a built-in bypass mechanism. For all other fiber-optic connections on the FortiDDoS 2000B and on other FortiDDoS models, it is recommended that you use FortiBridge bypass switches for failover protection.

For other bypass switches available in the market, please contact your Sales Engineer to check if the switch can work with FortiDDoS appliances. For more information, see “External bypass switches for maintenance & failover” on page 55.

Basic web hosting deployment

More complex setups can protect multiple systems. In a basic web hosting deployment, a FortiDDoS appliance can protect multiple customer systems as shown in Figure 15. You can use a system with a single Service Protection Profile (SPP) or multiple profiles. See “Key concepts” on page 16 for concepts related to profiles and “Identifying IP addresses and subnets to protect (SPP creation)” on page 111.

Figure 15: Basic web hosting deployment of FortiDDoS appliances

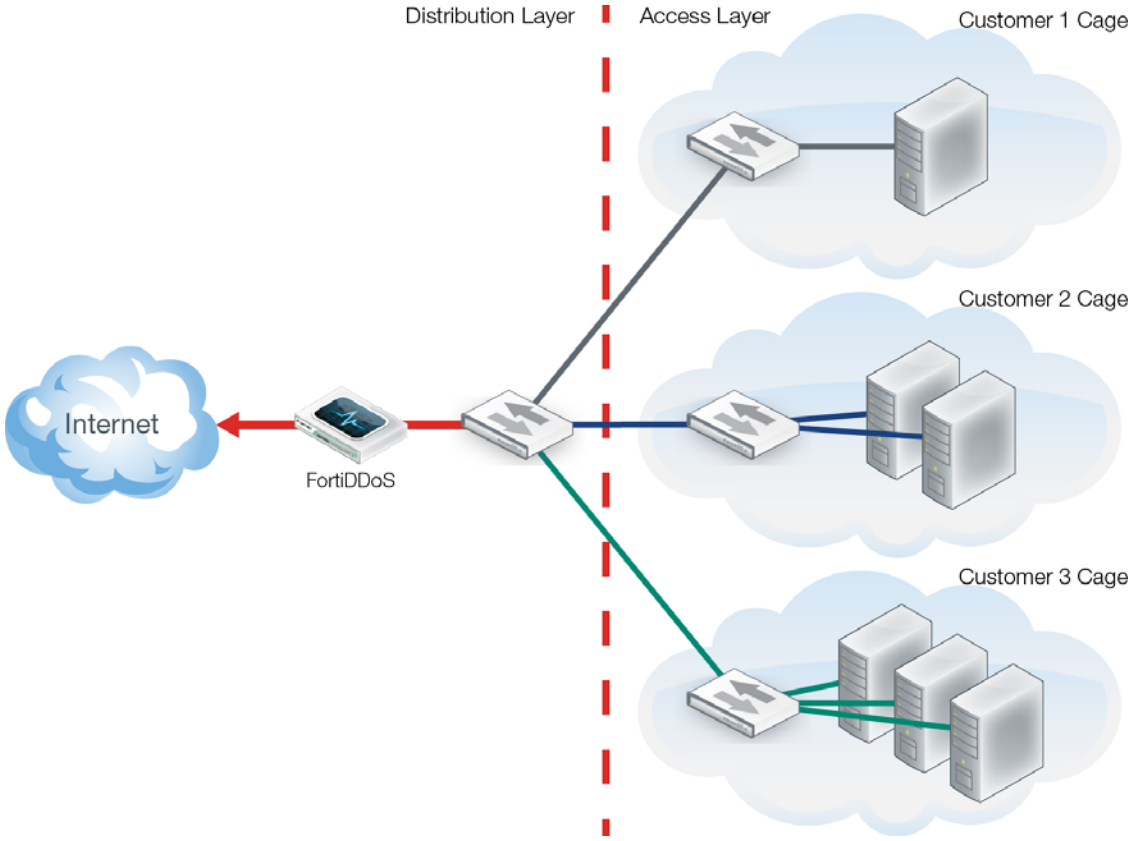

External bypass switches for maintenance & failover

The following FortiDDoS network interface connections have a built-in bypass mechanism:

- Any copper (RJ-45) network connections (for example, the RJ-45 connections for ports 1-16 on FortiDDoS 400 or 800)
- Ports 17-20 on the FortiDDoS 2000B, which are small form-factor pluggable (SFP; fiber-optic) links

In bypass mode, the FortiDDoS appliance automatically passes traffic through without performing any monitoring or prevention tasks. Packets that arrive at ingress ports are simply transferred to the corresponding egress ports, just like a wire.

FortiDDoS activates this bypass feature under the following conditions:

- The appliance is not powered up or is starting up or rebooting
- The appliance's FortiASIC processor or integrated switch fabric fail

This automatic bypass functionality is not available for the other fiber-optic connections on the FortiDDoS 2000B (ports 1-17) or for any of the fiber-optic connections found on other models.

For fiber-optic links that do not have a built-in bypass mechanism, you can use bypass switches to maintain connectivity. When both the FortiDDoS appliance and the failover switch share the same power supply, external connectivity is maintained during a power failure.

As shown in Figure 16, when the bypass switch is in disabled mode (its default mode), the inline traffic flows through the FortiDDoS appliance.

Figure 16: FortiDDoS in inline mode (switch with bypass mode disabled)

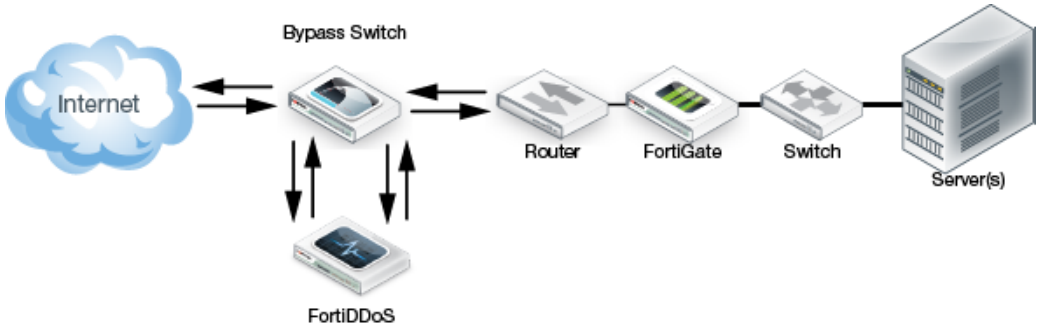

In Figure 17, the optical bypass switch is in bypass-enabled mode and all inline traffic is routed through the switch. After power is restored to the bypass switch, network traffic is diverted to the FortiDDoS appliance, allowing it to resume its critical functions.

Figure 17: FortiDDoS in bypass mode (switch with bypass mode enabled)

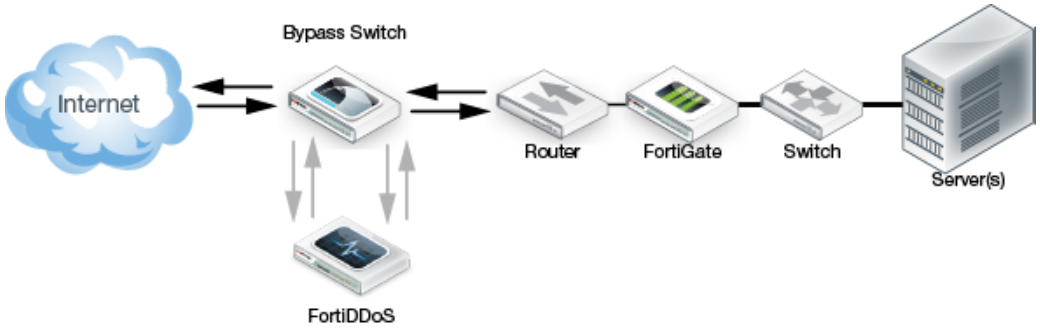

Either the automatic bypass mechanism or a bypass switch can maintain data traffic when there is a power or appliance failure. However, it is recommended that you automate failover behavior using a bypass switch with heartbeat. A bypass switch with heartbeat detects the failure of the FortiDDoS appliance (and the failure of traffic monitoring and mitigation) even when the appliance maintains the copper-based data link.

For information about detecting failures using the appliance's simple network management protocol (SNMP) agent, see "SNMP traps & queries" on page 243.

See also

- Using an optical bypass switch with heartbeat
- Configuring the optical bypass switch
- Connecting the optical bypass switch to the network and FortiDDoS
- Configuring MAC addresses for bypass switch heartbeat packets

Using an optical bypass switch with heartbeat

Fortinet recommends using a FortiBridge device as your optical bypass switch. This optical bypass switch shown in Figure 18 monitors the link to the attached FortiDDoS appliance by sending a heartbeat packet to the appliance once every second. If the optical bypass switch does not receive the heartbeat back, it automatically switches network traffic to bypass the unresponsive FortiDDoS appliance, even if the appliance is still receiving power. The optical bypass continues to send the heartbeat and restores the traffic through the FortiDDoS appliance as soon as the link is restored.

Figure 18: Optical bypass device

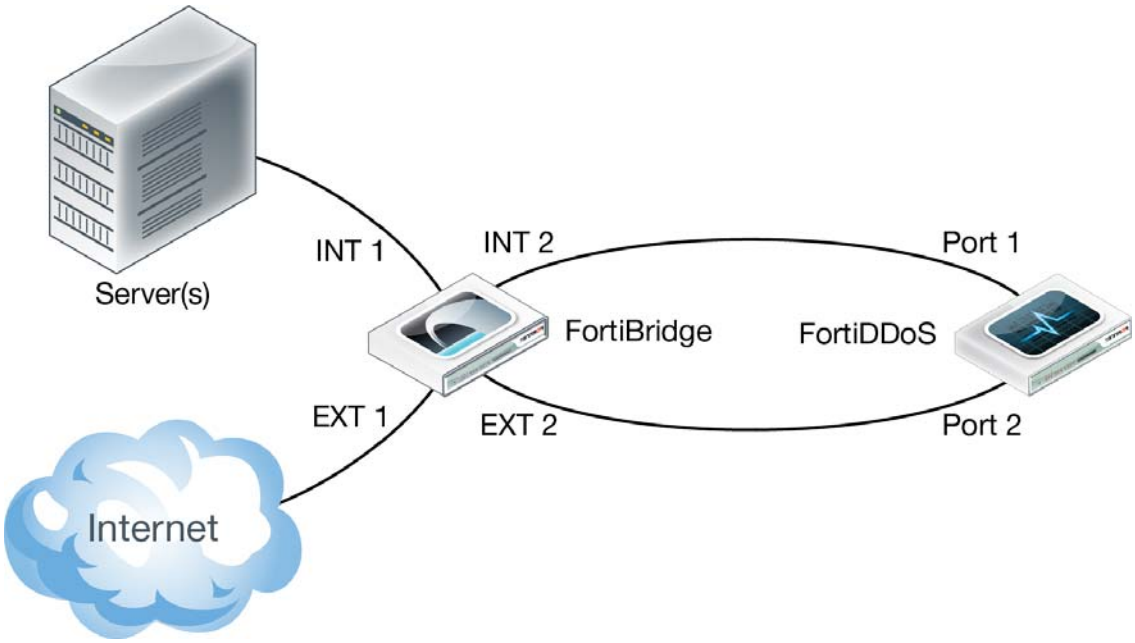

Configuring the optical bypass switch

Refer to the *FortiBridge QuickStart Guide* and *FortiGate Hardware Guide* to set the following parameters:

- Input timeout period
- Input retry count

Connecting the optical bypass switch to the network and FortiDDoS

1. Connect the INT 1 port to the Ethernet segment.
2. Connect the EXT 1 port to the Internet side.
3. Connect the INT 2 port to the FortiDDoS server port (for example, Port 1).
4. Connect the EXT 2 port to the FortiDDoS Internet port (for example, Port 2).

Configuring MAC addresses for bypass switch heartbeat packets

When a FortiDDoS appliance is used in conjunction with a bypass switch such as FortiBridge, ensure that FortiDDoS allows heartbeat packets from the bypass switch in all possible cases.

Typical bypass switches use heartbeat packets to check if the data path is connected. If the data path is broken for some critical reason, the bypass switch switches to bypass mode from normal mode.

To ensure that it passes on the heartbeat packets, FortiDDoS allows you to specify the MAC addresses that the bypass switch uses for the packets.

You can view these MAC addresses in the FortiBridge status page.

Every FortiDDoS link pair can be connected via a FortiBridge link pair. For example, you can use a FortiBridge link to bridge the Port 1/Port 2 link pair and another FortiBridge link to bridge the Port 3/Port 4 link pair. Each link pair is associated with a pair of MAC addresses. Therefore, if you are using two links, you configure four MAC addresses. If you are using one link, specify two MAC addresses.

You can program up to 16 MAC addresses.

If the bypass switches are from the same vendor, the most significant 24-bits of their MAC addresses are the same.

1. Click *Global Settings* > *Bypass MAC* > *Bypass MAC*.
2. Click *Add*, and then enter a name for the MAC address and the address.
3. Click *Save*.

Load balancing

Many data center and server farm architectures require network infrastructure to protect them. However, traffic volumes on some networks can exceed the capabilities of a single link pair on a FortiDDoS appliance or even the maximum throughput of a single appliance.

Table 5: Maximum throughput by model

Model	Maximum throughput (full duplex)	
	Per port pair	Per appliance
400B	1 Gbps	4 Gbps
800B	1 Gbps	8 Gbps
1000B	10 Gbps	12 Gbps
2000B	10 Gbps	24 Gbps

To increase the overall throughput, some topologies require some type of load-balancing solution using multiple link pairs or multiple FortiDDoS appliances.

The capacity of the load-balancing device must exceed the combined throughput of the multiple FortiDDoS appliances.

The load-balancing device intercepts all traffic between the server side and the Internet side and dynamically distributes the load among the available FortiDDoS appliances, based on the device's configuration. Load balancing utilizes all the appliances concurrently, providing overall improved performance, scalability and availability.

The FortiDDoS appliance is a layer-2 bridge and therefore does not have either a MAC address or an IP address in the data path. For transparent bridges, the load-balancing device receives a packet, makes a load-balancing decision, and forwards the packet to a FortiDDoS appliance. The FortiDDoS appliance does not perform NAT on the packets; the source and destination IP addresses are not changed.

The load-balancing device performs the following tasks:

- Balances traffic across two or more FortiDDoS appliances in your network, allowing them to work in parallel.
- Maintains state information about the traffic that flows through it and ensures that all traffic between specific IP address source and destination pairs flows through the same FortiDDoS appliance.
- Performs health checks on all paths through the FortiDDoS appliances. If any path is not operational, the load balancer maintains connectivity by diverting traffic away from that path.

You can use an external load balancer such as Linux Virtual Server (LVS), Cisco Content Switching Module (CSM), or Avaya Load Balancing Manager.

Load Balancing allows you to:

- Maximize FortiDDoS productivity
- Scale FortiDDoS performance
- Eliminate the FortiDDoS appliance as a single point of failure

Load balancing for FortiDDoS appliances requires a sandwich topology.

See also

- Sandwich topology for load balancing
- Switch configuration for load balancing

Sandwich topology for load balancing

The sandwich topology shown in Figure 19 places a load-balancing device before and after a pair of FortiDDoS appliances. For example, two 400B appliances to support a total throughput of 8 Gbps. This same topology and throughput is possible using a single 800B appliance.

This type of design ensures the highest level of security because it physically separates the FortiDDoS interfaces using multiple switches.

Each load-balancing device balances traffic between IP address interfaces of the peer device behind the FortiDDoS appliance. Each FortiDDoS appliance resides in a different VLAN and subnet and the physical ports connected to the FortiDDoS appliance are also on different VLANs. In addition, for each VLAN, both load-balancing devices are in the same subnet. Each load balancer interface and the FortiDDoS appliance connected to it reside in a separate VLAN. This configuration ensures persistency because all the traffic through a particular FortiDDoS appliance is contained in the appliance's VLAN.

In a typical load-balancing device, there are two hash predictors:

- **Bidirectional hash** requires both load-balancing devices to share a common hash value that ultimately produces the same route. You create bidirectional hashing by hashing the source and destination IP address along with the destination port of the given flow. The load-balancing devices ensure that all packets belonging to a session pass through the same FortiDDoS appliance in both directions. The devices select a FortiDDoS appliance based on a symmetric hash function of the source and destination IP addresses. This

ensures that packets traveling between the same source and destination IP addresses traverse the same FortiDDoS appliance.

- **Unidirectional hash** produces the route in the same fashion as a bidirectional hash and also creates a TCP connection table with the reverse flow path defined. This allows you to match return path traffic against this connection table rather than being hashed.

Figure 19: Load balancing using FortiDDoS appliances using sandwich topology

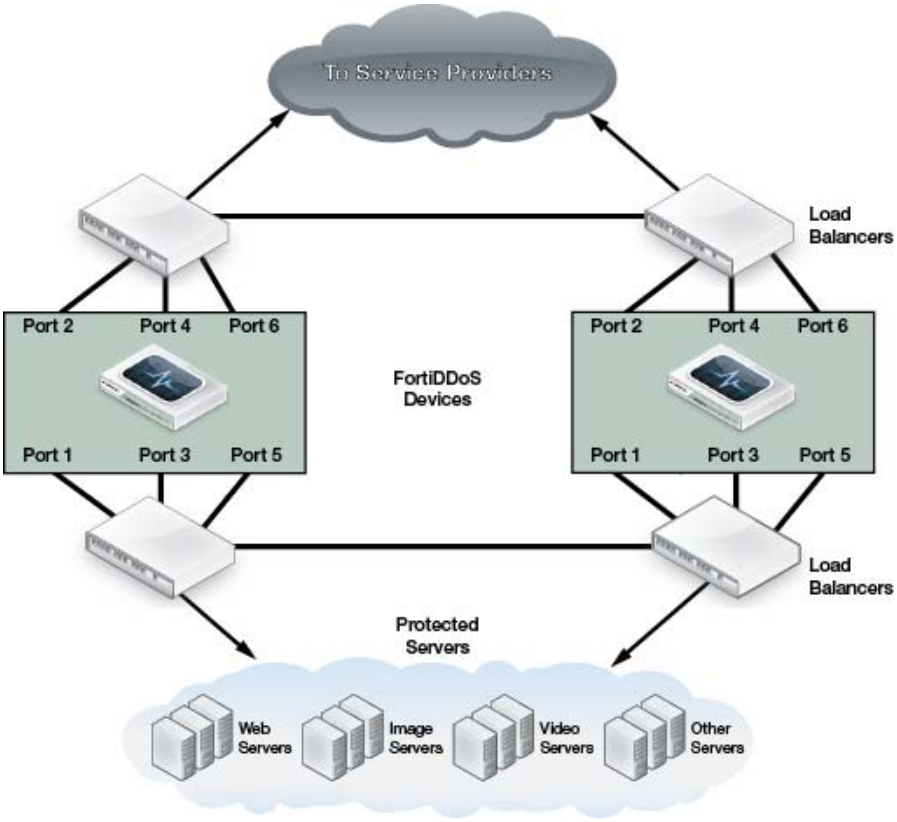

Switch configuration for load balancing using FortiSwitch

For an example configuration for the FortiSwitch 248-B DPS Ethernet switch, see “Appendix B: Switch & router configuration” on page 273.

Traffic diversion

In some environments, such as a service provider environment, the total bandwidth is more than what the FortiDDoS appliance supports. However, the attack traffic to a specific subnet or server is within the appliance’s capacity. You can route normal traffic through its regular path and manually divert the attack traffic. The FortiDDoS cleanses the diverted traffic and injects it back to the network.

The FortiDDoS appliance is a layer-2 bridge and therefore does not have either a MAC address or an IP address in the data path. To allow traffic to be diverted, connect the appliance to interfaces on the routers or switches that have a routeable IP address.

The example topology (shown in Figure 20) uses the following terminology:

- **Divert-from router:** Router from which the FortiDDoS appliance diverts the attacked customer traffic.
- **Inject-to router:** Router to which the FortiDDoS appliance forwards legitimate traffic.

See also

- Traffic diversion using separate divert-from and inject-to routers
- Traffic diversion using a single divert-from and inject-to router and a switch
- Setting thresholds for diverted traffic
- Configuring the routers & switch for traffic diversion

Traffic diversion using separate divert-from and inject-to routers

In the very simple deployment shown in Figure 20, Layer 2 forwards traffic through the FortiDDoS appliance.

An additional interface on Router 1 Divert-from Router diverts the traffic that is destined for the attacked destination. This traffic passes through the FortiDDoS appliance. The traffic is then forwarded to Router 2 Inject-to Router. These two interfaces are in the same network (192.168.1.x) and therefore an ARP request from Router 1 for 192.168.1.2 passes through the FortiDDoS appliance and reaches Router 2 and Router 2 can respond back with an ARP reply and vice versa.

A static route is added on Router 1 for addresses for the attacked customer network. Because it has the longest matching prefix, the rule matches first and therefore all traffic to the attacked customer network is diverted from Router 1 to Router 2 through the FortiDDoS appliance network rather than going straight from Router 1 to Router 2. Preferably, the return path for traffic is also through the FortiDDoS appliance. Although this solution works even if the traffic is unidirectional through the FortiDDoS appliance, bidirectional traffic helps the appliance determine the statefulness within connections.

Figure 20: Traffic diversion and a FortiDDoS appliance

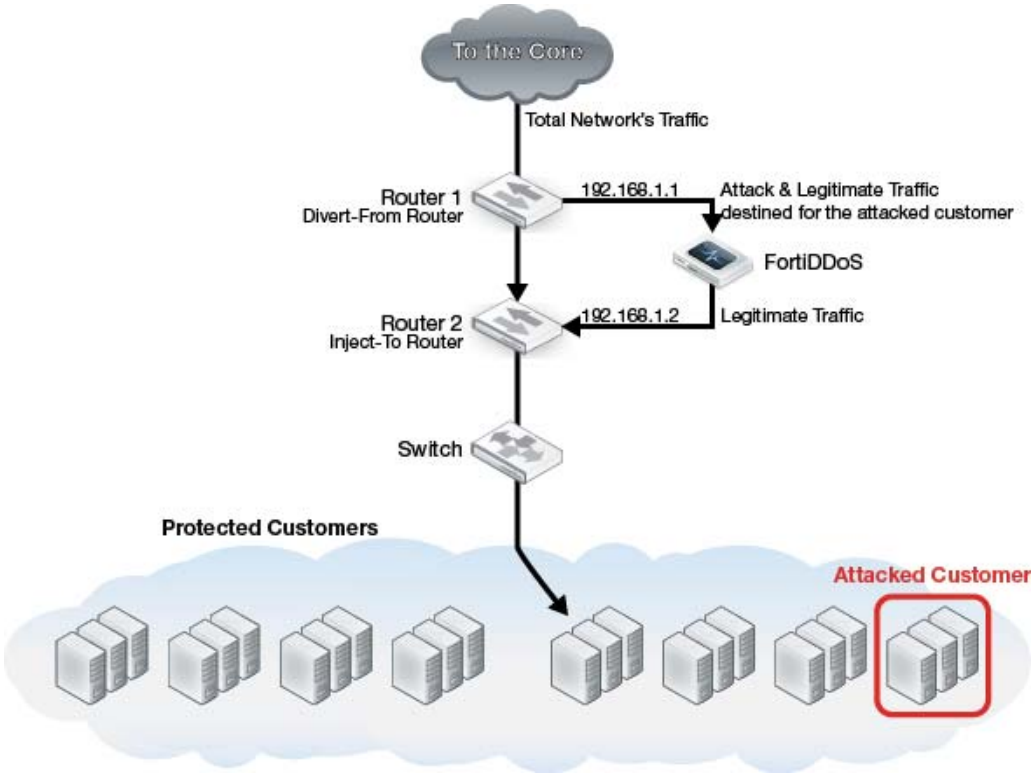

Traffic diversion using a single divert-from and inject-to router and a switch

Figure 21 shows a single router that is acting as both a divert-from and inject-to router. Layer 2 forwards through the FortiDDoS appliance.

One interface on the Internet side of the router diverts traffic to the attacked destination. This traffic passes through the FortiDDoS appliance through a switch. The traffic is then forwarded to the inject-to interface on the router through the same switch.

To ensure that the traffic is symmetric and both incoming and outgoing traffic to and from the attacked destination go through the FortiDDoS appliance, the LAN interface of the router diverts the traffic from the attacked destination. This traffic passes through the FortiDDoS appliance through a switch. The traffic is then forwarded to the inject-to interface on the same router through the same switch.

A static route is added on the router for addresses for the attacked customer network. Because it has the longest matching prefix, the rule matches first and therefore all traffic to the attacked customer network is diverted to the Layer 3 switch through the FortiDDoS appliance rather than going straight from the router to the distribution switch.

Preferably, the return path for traffic is through a FortiDDoS appliance. Although the solution works even if the traffic is unidirectional through the FortiDDoS appliance, bidirectional traffic helps the appliance determine the statefulness within connections.

To ensure that the return traffic passes through the FortiDDoS appliance, use the Policy Based Routing (PBR) that is available in most routers. PBR allows you to base routing on the source address of the packets and interface.

Figure 21: Traffic diversion using a single divert-from and inject-to router and a FortiDDoS appliance

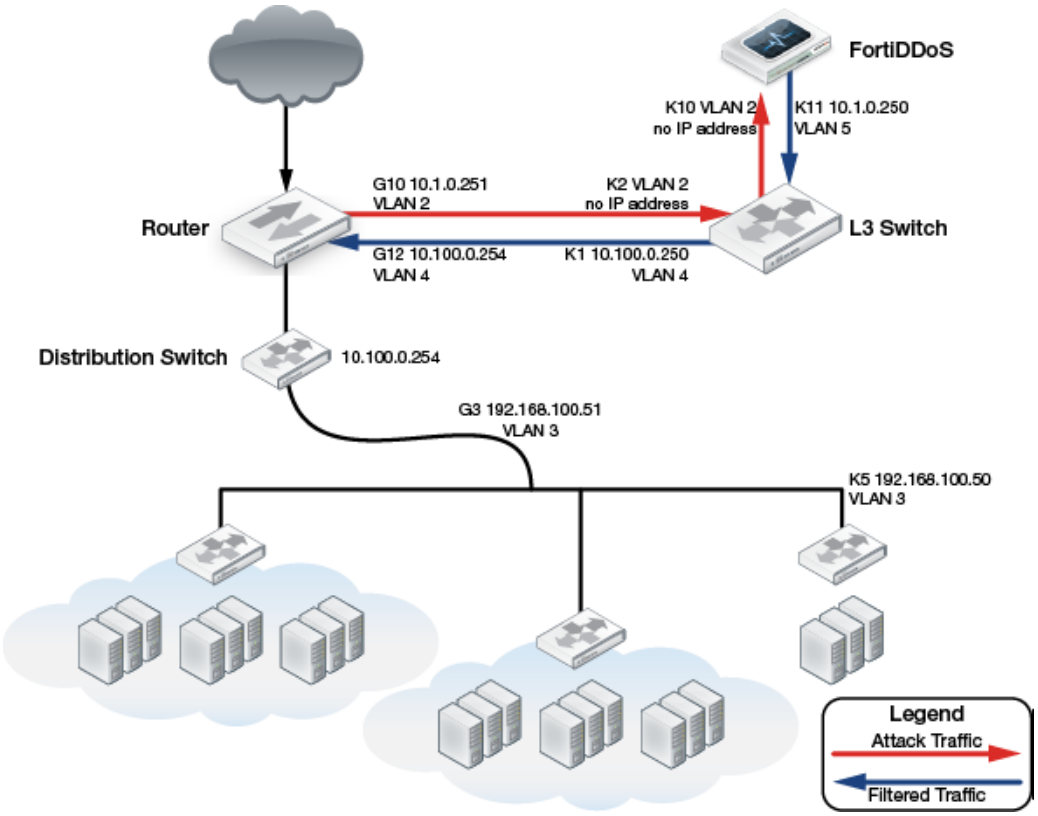

Router & switch configuration for diversion

For an example router and switch configuration for traffic diversion, see "Appendix B: Switch & router configuration" on page 273.

Setting thresholds for diverted traffic

In some cases, when traffic for a customer network is diverted through the FortiDDoS appliance during attacks, the appliance does not have traffic thresholds that correspond to the diverted network's normal traffic level or characteristics.

To solve this issue, do one of the following:

- **Archive a learning period:** During a time of normal traffic activity, divert the customer network traffic to a FortiDDoS appliance. Then, archive the configuration file created during this learning period using System > Maintenance > Backup & Restore. During an attack, restore the configuration and divert the affected traffic to the appliance with the restored configuration.
- **Create predefined profiles:** Create a series of backup configurations for different traffic levels. For example, define normal traffic levels for 1 Mbps, 10 Mbps, 20 Mbps, 100 Mbps, and so on. In addition, predefine profile parameters such as SYN/second, SYNs/Src, Concurrent Connections/Source, and so on. During an attack, restore the configuration that corresponds to the customer traffic level based on past historical knowledge of the data, and then divert the affected traffic to the appliance with the restored configuration.

  For information on setting key thresholds, see "Set using Emergency Setup" on page 154 and "Adjusting thresholds individually" on page 155.

Topology for synchronizing the configuration of two FortiDDoS appliances

The High Availability (HA) feature can synchronize configuration information between two FortiDDoS appliances to create an up-to-date standby appliance.

During events such as hardware failure or maintenance periods, the standby appliance can assume the role of the active appliance. You can manually change the routing to the standby appliance or use the failover configuration of switches or routers connected to the HA pair to automatically change it.

FortiDDoS can synchronize appliance configuration settings such as network interfaces, Service Protection Profile (SPP) and subnet definitions, and ACLs. For a list of settings and data that are not synchronized, see “Data and configuration settings that are not synchronized by HA”.

The synchronized configuration requires the following resources:

- Two identical physical FortiDDoS appliances

  Two appliances that are the same hardware model and have the same firmware version. For example, two FortiDDoS 400B appliances running firmware version 4.1.
- Redundant network topology

  Physical network cabling and routes that allow you to redirect web traffic to the standby appliance if the active appliance fails.

In the network topology shown in Figure 22, a single FortiDDoS appliance has been replaced with two appliances operating as an ***active-passive*** high availability (HA) pair. If the active appliance fails, the standby appliance is ready to assume the role of the failed appliance using the same configuration as the primary appliance. The standby appliance does not contain the traffic history or estimated threshold information found on the primary appliance.

To carry heartbeat and synchronization traffic between the HA pair, the heartbeat interfaces on the appliances (MGMT 2 port) are connected through crossover cables or a switch.

Figure 22: Configuration synchronization between two FortiDDoS appliances

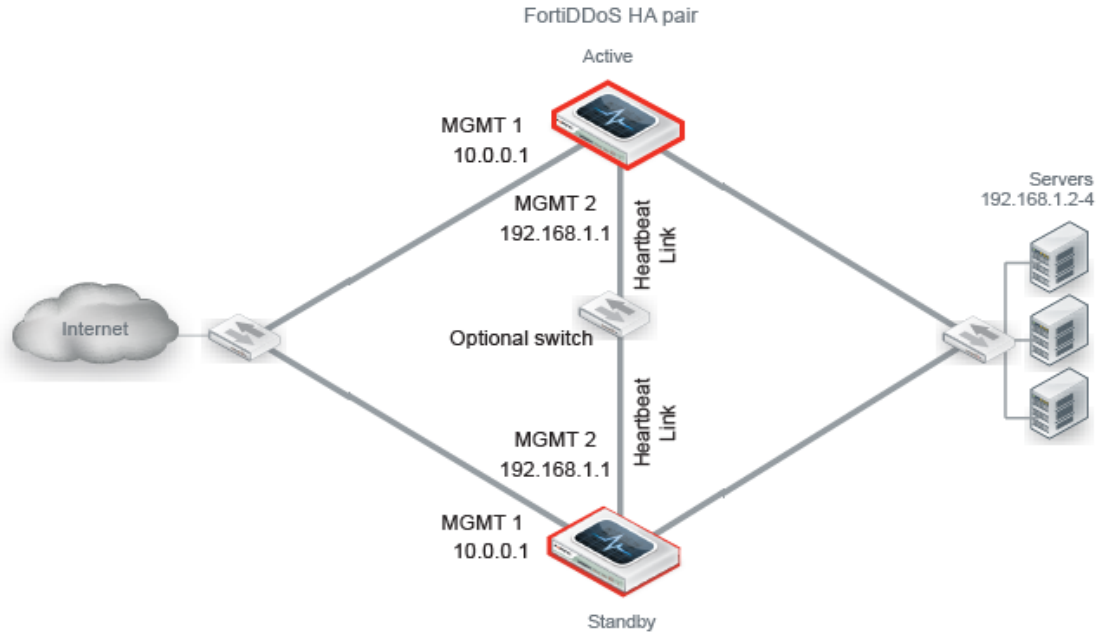

This configuration has the following limitations:

- If you use a switch to connect the heartbeat interfaces, they must be reachable by Layer 2 multicast.
- When the appliance fails, either an optical bypass switch or the FortiDDoS automatic bypass mechanism, which works for copper connections, can maintain the traffic flow. Because traffic is not interrupted, the standby appliance does not automatically assume the role of the failed appliance. In this case, use a manual process to make the standby appliance the active appliance.

  For more information about bypass configurations, see “External bypass switches for maintenance & failover” on page 55.

You can use SNMP to receive notification of the failure of a FortiDDoS device. For more information, see “SNMP traps & queries” on page 243.

When the former active appliance comes back online, it may or may not assume the active role. For an explanation, see “How HA chooses the active appliance” on page 67.

See also

- Heartbeat link and synchronization
- Data and configuration settings that are not synchronized by HA
- How HA chooses the active appliance
- Configuring configuration synchronization

Heartbeat link and synchronization

The heartbeat traffic indicates to the standby appliance in the HA pair that the main appliance is up and “alive.”

Heartbeat and synchronization traffic between paired appliances occur over the physical network port selected in *System* > *Config* > *High Availability*. HA traffic uses multicast UDP on

the MGMT 2 port. The HA multicast IP address is 239.0.0.1; it is hard-coded, and cannot be configured.

If switches are used to connect heartbeat interfaces between an HA pair, the heartbeat interfaces must be reachable by Layer 2 multicast.

To keep the configuration of the standby appliance up-to-date, HA pairs use the heartbeat link to automatically synchronize most of their configuration. Synchronization includes:

- core CLI-style configuration file (fddos_system.conf)
- X.509 certificates, certificate request files (CSR), and private keys

Synchronization occurs as soon as an appliance joins the pair and every 30 seconds thereafter. For a list of settings and data that are not synchronized, see "Data and configuration settings that are not synchronized by HA".

If a switch or router connects the heartbeat link, and the active appliance is very busy, it might require more time to establish the heartbeat link that determines the active appliance. You can configure the amount of time that a FortiDDoS appliance waits after it boots to establish this connection before assuming that the other appliance is unresponsive.

Data and configuration settings that are not synchronized by HA

FortiDDoS synchronizes appliance configuration settings such as network interfaces, Service Protection Profile (SPP) definitions, and ACLs. It does not synchronize session information or any other element of the data traffic.

It also does not synchronize recommended and estimated threshold information.

Configuration settings

All configuration settings on the active appliance are copied to the standby appliance, except the following settings:

Setting	Explanation
Host name	The host name is unique for each member of the FortiDDoS HA pair. See "Changing the FortiDDoS appliance's host name" on page 184.
HA active/standby status and priority	The *Device Priority* value is different on the main and standby appliances.

Log messages and generated reports

- **Log messages —** Describe system events that happened on a specific appliance. For example, log messages generated when the standby is acting as the active appliance (if you have configured local log storage) are stored on the standby appliance. For more information on configuring local log storage, see "Configuring system event logging" on page 235.
- **Generated reports —** Like the log messages that they are based upon, PDF, HTML, RTF, and plain text reports also describe events that happened on a specific appliance. Report settings are synchronized, but report output is not. For information about this feature, see "Reports" on page 250.

How HA chooses the active appliance

An HA pair may or may not resume their active and standby roles after the failed appliance starts to respond to the heartbeat again.

The current active appliance has, by definition, been available and working for a longer length of time (longer uptime) than a failed previous active appliance that has just returned online. Because it has the longer uptime, the current active appliance usually retains its status as the active appliance. However, if *Override* is enabled and the *Device Priority* setting of the returning appliance is higher, it is selected as the active appliance in the HA pair.

When *Override* is disabled, FortiDDoS determines the active appliance using the following criteria (in order):

1. The highest uptime value

   Uptime is reset to zero when an appliance fails or the status of any monitored port (specified by *Monitor/ Heartbeat*) changes.
2. The lowest *Device Priority* number (1 is the highest priority)
3. The highest-sorting serial number

   Serial numbers are sorted by comparing each character from left to right, where 9 and z are the greatest values, and result in highest placement in the sorted list.

Configuring configuration synchronization

To configure FortiDDoS appliances that are operating in HA mode, you usually connect only to the active (main) appliance. The active appliance's configuration is almost entirely synchronized to the passive appliance to ensure that it is prepared for failover routing.

However, you connect to the standby appliance for the following tasks:

- Viewing system log messages it records about itself on its hard disk.
- Configuring settings that are not synchronized. See "Data and configuration settings that are not synchronized by HA" on page 66.

Because the MGMT 1 network interface IP address of the standby appliance is synchronized with the IP address for the active appliance, you cannot connect to the web UI or CLI using the MGMT 1 port. To configure the standby appliance after you have added it to the HA pair, use a local console connection (see "Connecting to the CLI" on page 77).

To configure configuration synchronization (HA) via the web UI

1. If the HA pair uses FortiGuard IP Reputation Service, license both FortiDDoS appliances and register them with the Fortinet Technical Support web site:

   https://support.fortinet.com/

   For more information, see "FortiGuard IP Reputation Service" on page 130.
2. Cable both appliances into a redundant network topology.

   For an example, see Figure 22 on page 65.

3. Physically link the FortiDDoS appliances.

   To link the MGMT 2/HA ports for heartbeat and synchronization traffic between members of the pair, do one of the following:

   - Use a crossover cable to link the two appliances directly.
   - Link the appliances through a switch.

   If a switch is used to connect the heartbeat interfaces, the heartbeat interfaces must be reachable by Layer 2 multicast.

Maintain the heartbeat link(s). If the heartbeat is accidentally interrupted for an active-passive HA group, such as when a network cable is temporarily disconnected, the standby appliance assumes that the main appliance has failed. If no failure has actually occurred, both FortiDDoS appliances operate as primary appliances simultaneously. However, no traffic is routed through the standby device unless you use SNMP traps or polices on a switch or router to automatically re-route traffic.

If you link HA appliances through switches, to improve fault tolerance and reliability, link the ports through two separate switches. Do not connect these switches to your overall network, which could introduce a potential attack point. Connections to the overall network could also allow network load to cause latency in the heartbeat, which could cause an unintentional failover.

4. Log in to both appliances as the `admin` administrator account.

   To configure HA, an account requires an access profile includes *Read-Write* permissions to the *System* area. However, these accounts may not be able to use all the features that you use with HA, such as logs and network configuration. For details, see “Permissions” on page 44.

5. On both appliances, go to *System* > *Config* > *High Availability*.

6. For *Configured HA Mode*, select *Active-Passive*, and then configure the following settings:

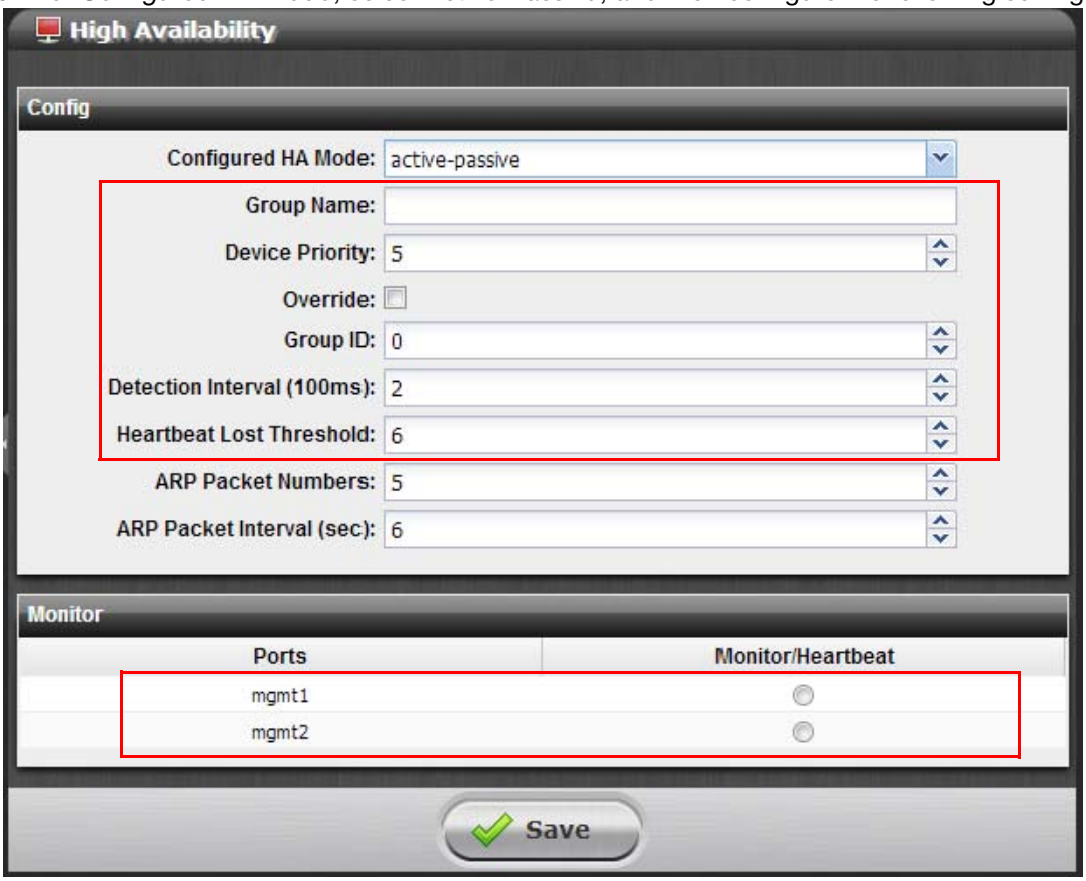

Setting name	Description
Group Name	Enter a name that identifies the HA pair, if you have more than one. This setting is optional, and does not affect HA function. The maximum length is 35 characters.
Device Priority	Type the priority of the appliance when selecting the active appliance in the HA pair. This setting is optional. The smaller the number, the higher the priority. The valid range is 0 to 9. The default is 5. **Note:** By default, unless you enable *Override*, the amount of time a device is working and available (uptime) is more important than this setting. For details, see "How HA chooses the active appliance" on page 67.
Override	Specify whether *Device Priority* has priority over appliance uptime when FortiDDoS selects the active appliance. See "How HA chooses the active appliance" on page 67.
Group ID	Type a number that identifies the HA pair. Both members of the HA pair must have the same group ID. If you have more than one HA pair on the same network, each HA pair must have a unique group ID. Changing the group ID changes the pair's virtual MAC address. The valid range is 0 to 63. The default value is 0.

Setting name	Description
Detection Interval	Specify the length of the pause between each heartbeat packet that the one FortiDDoS appliance sends to the other FortiDDoS appliance in the HA pair, in 100-millisecond intervals. This is also the amount of time that a FortiDDoS appliance waits before expecting to receive a heartbeat packet from the other appliance. This part of the configuration is synchronized between the active and standby appliances. The valid range is 1 to 20 (that is, between 100 and 2,000 milliseconds). **Note:** Although this setting is synchronized between the main and standby appliances, configure both appliances with the same interval to prevent inadvertent failover from occurring before the initial synchronization.
Heartbeat Lost Threshold	Type the number of times that the HA appliance retries the heartbeat and waits to receive HA heartbeat packets from the other HA appliance before it assumes that the other appliance has failed. This part of the configuration is synchronized between the main appliance and standby appliance. Normally, you do not change this setting. Exceptions include: • Increase the failure detection threshold when a failure is detected when none has actually occurred. For example, during peak traffic times when the main appliance is very busy and does not respond to heartbeat packets in time. The valid range is from 1 to 60. **Note:** Although this setting is synchronized between the main and standby appliances, configure both appliances with the same threshold to prevent inadvertent failover from occurring before the initial synchronization.
Monitor/ Heartbeat	Select a network interface for port monitoring and for sending heartbeat signals and synchronization data between the main and standby appliances. Port monitoring (also called interface monitoring) monitors physical network ports to verify that they are functioning properly and linked to their networks. If the physical port fails or the cable becomes disconnected, failover occurs. **Note:** To prevent an unintentional failover, do not configure port monitoring until after you configure HA on both appliances in the HA pair and have plugged in the cables to link the physical network ports that are monitored. In most cases, MGMT 2 is the port monitoring and heartbeat interface. **Note:** If a switch is used to connect the heartbeat interfaces, the heartbeat interfaces must be reachable by Layer 2 multicast.

7. Click *Save*.

   The appliances join the HA pair by matching their *Group ID*. They begin to send heartbeat and synchronization traffic to each other through their heartbeat links.

   To determine which appliance currently has the role of the main (active) appliance, on *System* > *Status* > *Dashboard*, in the System Information widget, see *Effective HA Role*.

   - *standalone* — The appliance is not configured as a member of an HA pair.
   - *main* — The configuration of the main appliance is synchronized with the configuration of the standby appliance. Also called the primary or master appliance.
   - *standby* — The configuration of the standby appliance is synchronized with the main appliance. The passive appliance listens to heartbeat traffic and port monitoring for signs that the main appliance has become unresponsive, at which point it maintains its configuration independently. Also called the secondary or slave appliance.

   If both appliances believe that they are the main:

   - Test the cables and switches in the heartbeat link to verify that the link is functional.
   - Verify that you have selected the heartbeat port in *Monitor/ Heartbeat*.
   - Verify that the appliances have the same *Group ID* value.
   - Verify that the ports specified by *Monitor/ Heartbeat* are linked and available.
   - For debugging logs, use the `diagnose debug application hasyncd level` command.

8. To monitor the HA pair for failover, you can use SNMP (see "Configuring SNMP settings for system alarms and event messages" on page 243), log messages, and alert email (see "Configuring system event logging" on page 235).

   If failover time is too long, adjust the following:

Setting name	Description
ARP Packet Numbers	Type the number of times that the FortiDDoS appliance broadcasts extra address resolution protocol (ARP) packets when it takes on the main role. (Even though a new NIC has not actually been connected to the network, FortiDDoS does this to notify the network that a new physical port has become associated with the IP address and virtual MAC of the HA pair.) This is sometimes called "using gratuitous ARP packets to train the network," and can occur when the main appliance is starting up or during failover. Also configure *ARP Packet Interval*. Normally, you do not need to change this setting. If your HA pair takes a long time to train the network, you can increase the number of times the main appliance sends gratuitous ARP packets. Sending more gratuitous ARP packets helps the training process to happen faster. The valid range is 1 to 16.
ARP Packet Interval	Type the number of seconds to wait between each broadcast of ARP packets. Normally, you do not need to change this setting. If your HA pair takes a long time to train the network, you can decrease the interval. Sending ARP packets more frequently may help the training process to happen faster. The valid range is from 1 to 20.

To configure configuration synchronization via the CLI

Enter the following commands:

```
config system ha
  set mode <standalone | active-passive>
  set group-name <group_name_str>
  set priority <priority_int>
  set override <enable | disable>
  set group-id <group_id_integer>
  set hb-interval <hb_interval_int>
  set hb-lost-threshold <hb_lost_thresh_int>
  set hbdev <mgmt1 | mgmt2>
  set arps <arps_int>
  set arps-interval <arps_interval_int>
end
exit
```

where:

- `<standalone | active-passive>` specifies whether the appliance operates as one of a pair of appliances that synchronizes their configuration (HA pair).
- `<group_name_str>` specifies a name that identifies the HA pair.
- `<priority_int>` specifies the priority of the appliance when selecting the active appliance in the HA pair.
- `<enable | disable>` specifies whether priority has priority over appliance uptime when FortiDDoS selects the active appliance.
- `<group_id_integer>` specifies a number between 0 and 63 that identifies the members of an HA pair.
- `<hb_interval_int>` specifies the length of the pause between each heartbeat packet that the one FortiDDoS appliance sends to the other FortiDDoS appliance in the HA pair, in 100-millisecond intervals.
- `<hb_lost_thresh_int>` specifies the number of times that the HA appliance retries the heartbeat before it assumes that the other HA appliance has failed.
- `<mgmt1 | mgmt2>` specifies a network interface for port monitoring and for sending heartbeat signals and synchronization data between the main and standby appliances.
- `<monitor_str>` specifies
- `<arps_int>` specifies the number of times that the FortiDDoS appliance broadcasts extra address resolution protocol (ARP) packets when it takes on the main role. The valid range is from 1 to 16.
- `<arps_interval_int>` specifies the number of seconds to wait between each broadcast of ARP packets. The valid range is from 1 to 20.

For more detailed information about the settings, see “To configure configuration synchronization (HA) via the web UI” on page 67.

Feature availability in a one-way traffic configuration

You can configure FortiDDoS to handle traffic that is asymmetric — only incoming or outgoing traffic travels through the appliance. For example, you can route only inbound traffic through FortiDDoS.

Almost all of the FortiDDoS attack mitigation features work with an asymmetric traffic flow, including all layer 3 and layer 7 flood mitigation and mitigation for layer 4 floods such as TCP and UDP port floods, SYN floods, and zombie floods.

However, the following FortiDDoS features require a two-way traffic flow:

- TCP state violation flood mitigation
- TCP state and foreign packet anomaly recognition (all other types of anomaly recognition work with asymmetric traffic)

For more information on these features, see “TCP State Anomalies drop graph” on page 218.

In addition, when it handles an asymmetric traffic flow, FortiDDoS is unable to determine the precise number of concurrent connections. This limitation affects the following statistics:

- Connection Per Source (source flood)
- Connection Per Destination (destination flood)

For more information about these features, see “Layer 4 graphs” on page 220.

Connecting to the web UI or CLI

To configure, maintain, and administer the FortiDDoS appliance, you need to connect to it. There are two methods:

- **Web UI —** A graphical user interface (GUI), from within a web browser. It can display reports and logs but does not provide many advanced diagnostic commands. For usage, see "How to use the web UI" on page 43.

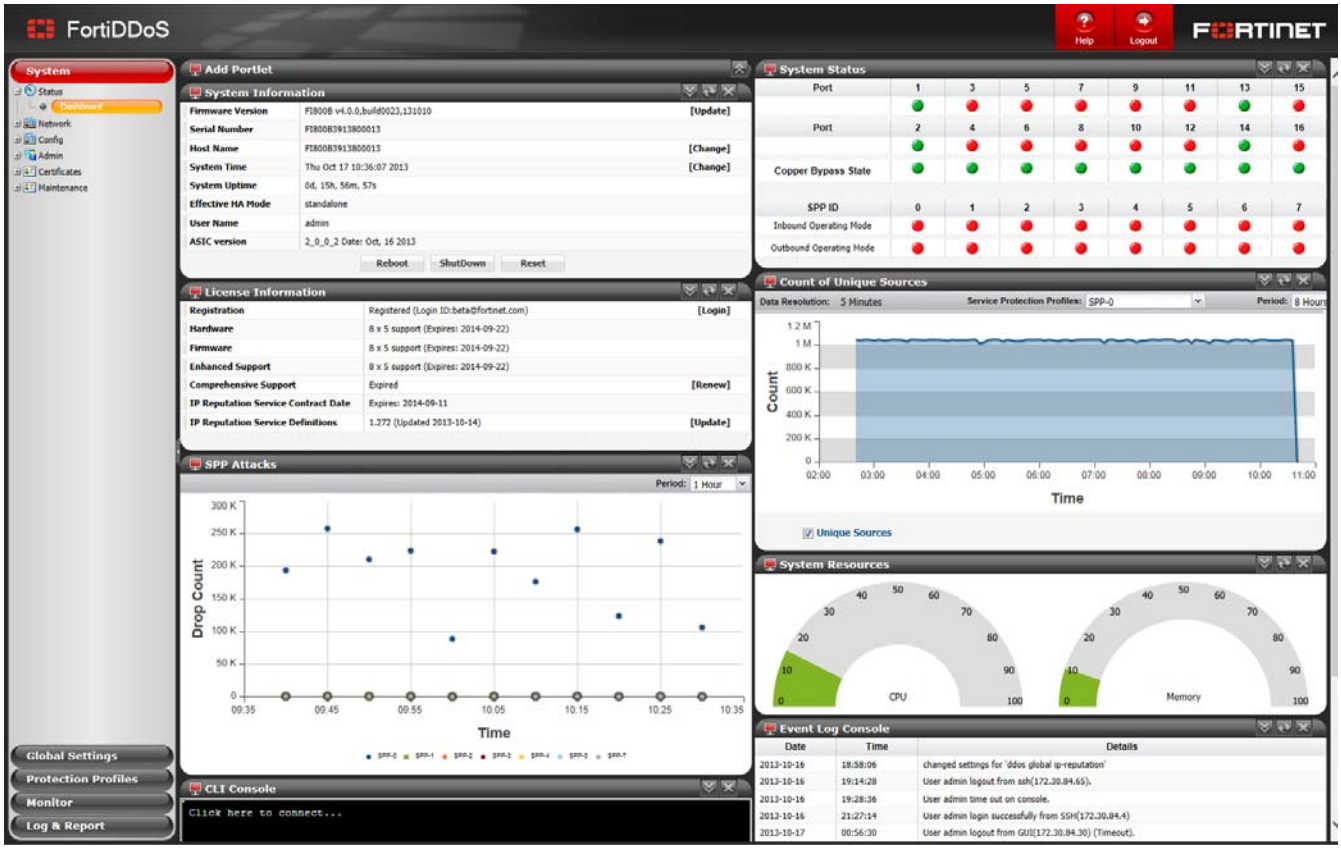

- **Command line interface (CLI) —** A text interface similar to DOS or UNIX commands, from a Secure Shell (SSH) or Telnet terminal, or from the JavaScript CLI Console widget in the

web UI (System > Status > Dashboard). It provides access to many advanced diagnostic commands as well as configuration, but lacks reports and logs.

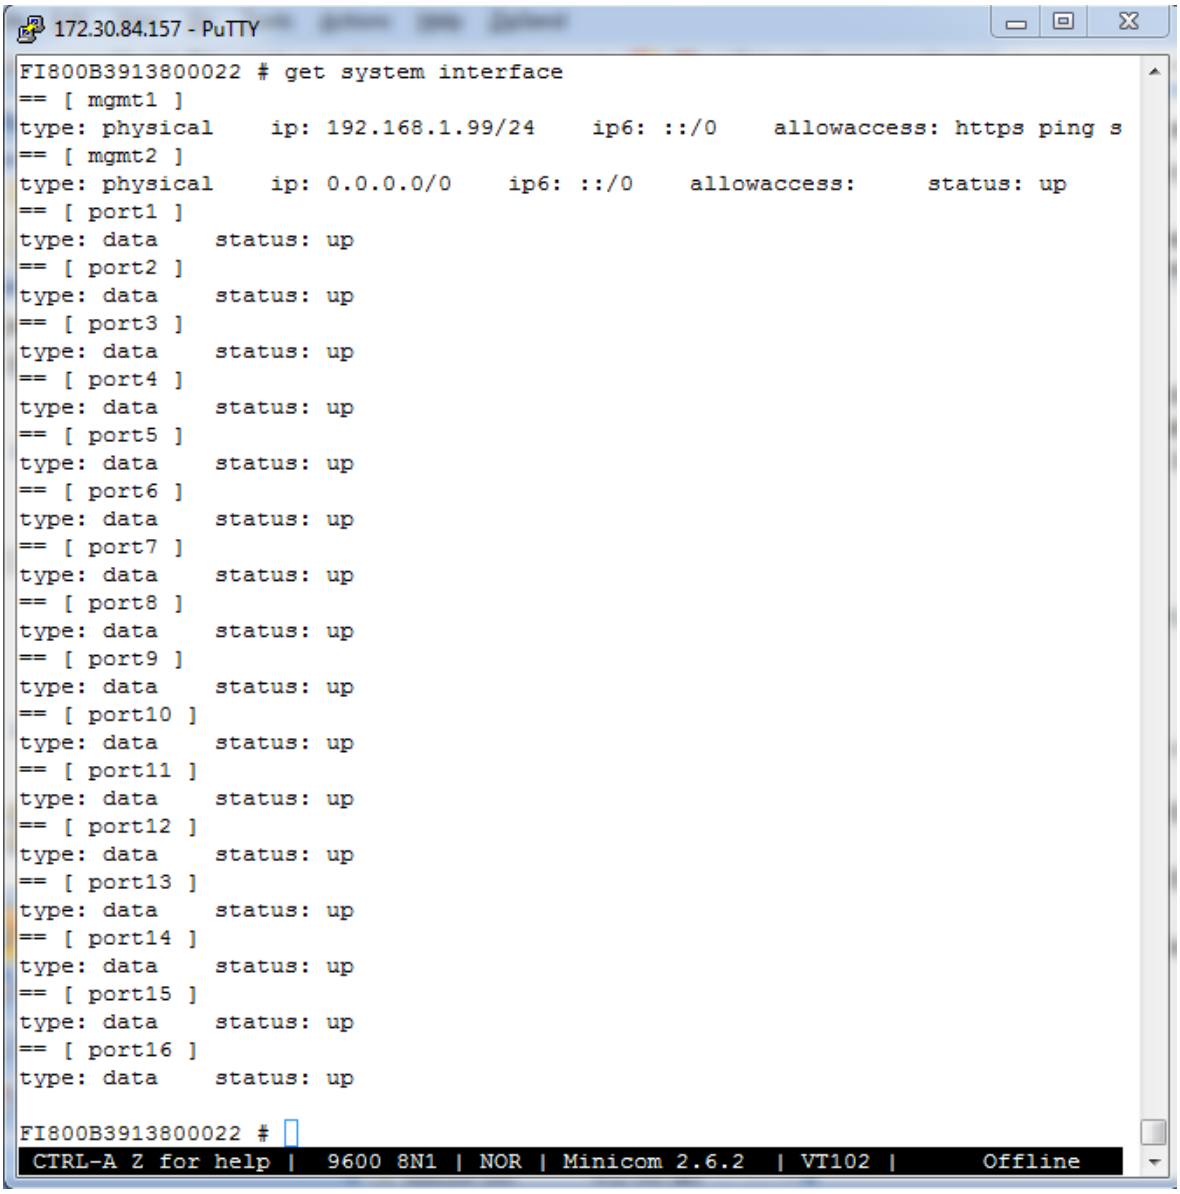

Access to the CLI and/or web UI through your network is not yet configured if:

- you are connecting for the first time
- you have just reset the configuration to its default state
- you have just restored the firmware

In these cases, you must initially connect your computer directly to FortiDDoS, using the default settings.

Via the direct connection, you can use the web UI or CLI to configure FortiDDoS's basic network settings. Once this is done, you can place FortiDDoS on your network, and use FortiDDoS through your network.

Until the FortiDDoS appliance is configured with an IP address and connected to your network, you may prefer to connect the FortiDDoS appliance directly to your management computer, or through a switch, in a peer network that is isolated from your overall network. This improves security during setup. However, isolation is not required.

Connecting to the web UI

You can connect to the web UI using its default settings.

Table 6: Default settings for connecting to the web UI

Network Interface	MGMT 1
URL	https://192.168.1.99/
Administrator Account	admin
Password	

Requirements

- a computer with an RJ-45 Ethernet network port
- a web browser such as Microsoft Internet Explorer version 8.0 or newer, or Mozilla Firefox 20 or newer
- a crossover Ethernet cable

To connect to the web UI

1. On your management computer, configure the Ethernet port with the static IP address 192.168.1.100 with a netmask of 255.255.255.0.
2. Using the Ethernet cable, connect your computer's Ethernet port to MGMT 1 on the FortiDDoS appliance.

3. Start your browser and enter the URL:

   https://192.168.1.99/

   (Remember to include the “s” in https://.)

   Your browser connects the appliance.

   To support HTTPS authentication, the FortiDDoS appliance ships with a self-signed security certificate, which it presents to clients whenever they initiate an HTTPS connection to the FortiDDoS appliance. When you connect, depending on your web browser and prior access of the FortiDDoS appliance, your browser might display two security warnings related to this certificate:

   - The certificate is not automatically trusted because it is self-signed, rather than being signed by a valid certificate authority (CA). Self-signed certificates cannot be verified with a proper CA, and therefore might be fraudulent. You must manually indicate whether or not to trust the certificate.
   - The certificate might belong to another web site. The common name (CN) field in the certificate, which usually contains the host name of the web site, does not exactly match the URL you requested. This could indicate server identity theft, but could also simply indicate that the certificate contains a domain name while you have entered an IP address. You must manually indicate whether this mismatch is normal or not.

     Because the CN field for the FortiDDoS certificate is the serial number of the appliance, you can use this field to verify you are using the correct appliance.

   Both warnings are normal for the default certificate. SSL v3 and TLS v1.0 are supported.

4. Verify and accept the certificate, either permanently (the web browser does not display the self-signing warning again) or temporarily. You cannot log in until you accept the certificate.

   For details on accepting the certificate, see the documentation for your web browser.

5. In the *Username* field, type `admin`, and then click *Login*. (In its default state, there is no password for this account.)

   Login credentials entered are encrypted before they are sent to the FortiDDoS appliance. If your login is successful, the web UI appears. To continue by updating the firmware, see “Updating the firmware” on page 81. Otherwise, to continue by setting an administrative password, see “Changing the “admin” account password” on page 93.

Connecting to the CLI

Using its default settings, you can access the CLI from your management computer in two ways:

- a local console connection
- an SSH connection, either local or through the network

Secure Shell (SSH) provides both secure authentication and secure communications to the CLI. Supported SSH protocol versions, ciphers, and bit strengths include SSH version 2 with AES-128, 3DES, Blowfish, and SHA-1.

Table 7: Default settings for connecting to the CLI by SSH

Network Interface	MGMT 1
IP Address	192.168.1.99
SSH Port Number	22
Administrator Account	admin
Password	

The default settings are in effect only when you connect for the first time, have reset the configuration to its default state, or restored the firmware. If administrative access settings have already been configured, access the CLI using the IP address, administrative access protocol, administrator account and password already configured, instead of the default settings.

Requirements

- a computer with an available serial communications (COM) port
- the RJ-45-to-DB-9 or null modem cable included in your FortiDDoS package
- terminal emulation software such as PuTTY or Tera Term

The following procedures describe connection using PuTTY software; steps may vary with other terminal emulators.

To connect to the CLI using a local console connection

1. Using the RJ-45-to-DB-9 or null modem cable, connect your computer's serial communications (COM) port to the FortiDDoS appliance's console port.
2. Verify that the FortiDDoS appliance is powered on.
3. On your management computer, start PuTTY.
4. In the *Category* tree on the left, go to *Connection* > *Serial* and configure the following:

Serial line to connect to	COM1 (or, if your computer has multiple serial ports, the name of the connected serial port)
Speed (baud)	9600
Data bits	8
Stop bits	1
Parity	None
Flow control	None

5. In the *Category* tree on the left, go to *Session* (***not*** the sub-node, *Logging*) and from *Connection type*, select *Serial*.
6. Click *Open*.

7. Press the Enter key to initiate a connection.

   The login prompt appears.

8. Type `admin` then press Enter twice. (In its default state, there is no password for the `admin` account.)

   The CLI displays a command line prompt. The hostname is the serial number of the appliance.

   You can now enter commands. To continue by updating the firmware, see "Updating the firmware" on page 81. Otherwise, to continue by setting an administrative password, see "Changing the "admin" account password" on page 93.

Requirements

- a computer with an RJ-45 Ethernet port
- a crossover Ethernet cable (if connecting directly) or straight-through Ethernet cable (if connecting through a switch or router)
- a FortiDDoS network interface configured to accept SSH connections (In its default state, Port 1 accepts SSH. You may need to connect directly first in order to configure a static route so that, later, you can connect through routers. For details, see "Adding a gateway" on page 103.)
- an SSH client, such as PuTTY

To connect to the CLI using an SSH connection

1. On your management computer, configure the Ethernet port with the static IP address 192.168.1.100 with a netmask of 255.255.255.0.
2. Using the Ethernet cable, connect your computer's Ethernet port to the console port on the FortiDDoS appliance.
3. Verify that the FortiDDoS appliance is powered on.
4. On your management computer, start PuTTY.

   Initially, the *Session* category of settings is displayed.

5. In *Host Name (or IP Address)*, type `192.168.1.99`.
6. In Port, type `22`.
7. From *Connection type*, select *SSH*.
8. Select *Open*.

   The SSH client connects to the FortiDDoS appliance.

   The SSH client may display a warning if this is the first time you are connecting to the FortiDDoS appliance and its SSH key is not yet recognized by your SSH client, or if you have previously connected to the FortiDDoS appliance but it used a different IP address or SSH key. If your management computer is directly connected to the FortiDDoS appliance with no network hosts between them, this is normal.

9. Click *Yes* to verify the fingerprint and accept the FortiDDoS appliance's SSH key. You cannot log in until you accept the key.

   The CLI displays a login prompt.

10. Type `admin` and press Enter. (In its default state, there is no password for this account.)

If three incorrect login or password attempts occur in a row, you are disconnected. Wait one minute, and then reconnect to attempt the login again.

The CLI displays a prompt with the serial number of the appliance. For example:

```
FI800B3913800022 #
```

You can now enter commands. To continue by updating the firmware, see “Updating the firmware” on page 81. Otherwise, to continue by setting an administrative password, see “Changing the “admin” account password” on page 93.

Updating the firmware

Your new FortiDDoS appliance comes with the latest operating system (firmware) when shipped. However, if a new version has been released since your appliance was shipped, you should install it before you continue the installation.

Fortinet periodically releases FortiDDoS firmware updates to include enhancements and address issues. After you register your FortiDDoS appliance, FortiDDoS firmware is available for download at:

https://support.fortinet.com

Installing new firmware overwrites any FortiGuard IP Reputation Service definitions and disables the service. After any firmware update, re-enable the IP Reputation feature. FortiDDoS downloads the current definitions as part of the enabling process. For more information, see “FortiGuard IP Reputation Service” on page 130.

New firmware can introduce new features which you must configure for the first time.

For late-breaking information specific to the firmware release version, see the Release Notes available with that release.

In addition to major releases that contain new features, Fortinet releases patch releases that resolve specific issues but do not contain new or changed features. It is recommended that you download and install patch releases as soon as they are available.

Before you can download firmware updates for your FortiDDoS appliance, you must first register your FortiDDoS appliance with Fortinet Technical Support. For details, go to https://support.fortinet.com/ or contact Fortinet Technical Support.

See also

- Testing new firmware before installing it
- Installing firmware
- Installing alternate firmware

Testing new firmware before installing it

You can test a new firmware image by temporarily running it from memory, without saving it to disk. By keeping your existing firmware on disk, if the evaluation fails, you do not have to re-install your previous firmware. Instead, you can quickly revert to your existing firmware by simply rebooting the FortiDDoS appliance.

To test a new firmware image

1. Download the firmware file from the Fortinet Technical Support web site:

   https://support.fortinet.com/
2. Connect your management computer to the FortiDDoS console port using a RJ-45-to-DB-9 serial cable or a null-modem cable.

3. Initiate a connection from your management computer to the CLI of the FortiDDoS appliance.

   For details, see "Connecting to the web UI or CLI" on page 74.
4. Connect MGMT 1 of the FortiDDoS appliance directly or to the same subnet as a TFTP server.
5. Copy the new firmware image file to the root directory of the TFTP server.
6. If necessary, start your TFTP server. (If you do not have one, you can temporarily install and run one such as `tftpd` (Windows, Mac OS X, or Linux) on your management computer.)

Because TFTP is ***not*** secure, and because it does not support authentication and could allow anyone to have read and write access, you should ***only*** run it on trusted administrator-only networks, ***never*** on computers directly connected to the Internet. If possible, immediately turn off `tftpd` off when you are done.

7. Verify that the TFTP server is currently running, and that the FortiDDoS appliance can reach the TFTP server.

   To use the FortiDDoS CLI to verify connectivity, enter the following command:

   ```
   execute ping 192.168.1.168
   ```

   where `192.168.1.168` is the IP address of the TFTP server.
8. Enter the following command to restart the FortiDDoS appliance:

   ```
   execute reboot
   ```
9. As the FortiDDoS appliances starts, a series of system startup messages appear.

   ```
   Press any key to display configuration menu.....
   ```
10. Immediately press a key to interrupt the system startup.

You have only three seconds to press a key. If you do not press a key soon enough, the FortiDDoS appliance reboots and you must log in and repeat the `execute reboot` command.

If you successfully interrupt the startup process, the following messages appear:

```
[G]:  Get firmware image from TFTP server.
[F]:  Format boot device.
[B]:  Boot with backup firmware and set as default.
[Q]:  Quit menu and continue to boot with default firmware.
[H]:  Display this list of options.

Enter G,F,B,Q,or H:

Please connect TFTP server to Ethernet port "MGMT".
```

11. Type `G` to get the firmware image from the TFTP server.

    The following message appears:

    ```
    Enter TFTP server address [192.168.1.168]:
    ```
12. Type the IP address of the TFTP server and press Enter.

    The following message appears:

    ```
    Enter local address [192.168.1.188]:
    ```

13. Type a temporary IP address that can be used by the FortiDDoS appliance to connect to the TFTP server.

    The following message appears:

    ```
    Enter firmware image file name [image.out]:
    ```

14. Type the firmware image file name and press Enter.

    The FortiDDoS appliance downloads the firmware image file from the TFTP server and displays a message similar to the following:

    ```
    MAC:00219B8F0D94
    ############################
    Total 28385179 bytes data downloaded.
    Verifying the integrity of the firmware image..
    Save as Default firmware/Backup firmware/Run image without
    saving:[D/B/R]?
    ```

If the download fails after the integrity check with the error message:

```
invalid compressed format (err=1)
```

but the firmware matches the integrity checksum on the Fortinet Technical Support web site, try a different TFTP server.

15. Type `R`.

    The FortiDDoS image is loaded into memory and uses the current configuration, ***without*** saving the new firmware image to disk.

16. To verify that the new firmware image was loaded, log in to the CLI and type:

    ```
    get system status
    ```

17. Test the new firmware image.

    - If the new firmware image operates successfully, you can install it to disk, overwriting the existing firmware, using the procedure "Installing firmware" on page 83.
    - If the new firmware image does ***not*** operate successfully, reboot the FortiDDoS appliance to discard the temporary firmware and resume operation using the existing firmware.

See also

- Installing firmware
- Installing alternate firmware

Installing firmware

You can use either the web UI or the CLI to upgrade or downgrade the appliance's operating system.

Firmware changes are either:

- an update to a newer version
- a reversion to an earlier version

To determine if you are updating or reverting the firmware, go to *System* > *Status* > *Dashboard* and in the *System Information* widget, see the *Firmware Version* row. (Alternatively, in the CLI, enter the command `get system status`.)

For example, if your current firmware version is:

`FI800B v4.0.0,build0011,130806`

changing to

`FI800B v4.0.0,build0010,130706`

an earlier build number (10) and date (`130706` means July 6, 2013), indicates that you are reverting.

Back up your configuration before beginning this procedure.

Reverting to an earlier firmware version could reset settings that are not compatible with the new firmware. For information on backups, see “Backups” on page 167. For information on reconnecting to a FortiDDoS appliance whose network interface configuration has been reset, see “Connecting to the web UI or CLI” on page 74.

To install firmware via the web UI

1. Download the firmware file from the Fortinet Technical Support web site:

   https://support.fortinet.com/

2. Log in to the web UI of the FortiDDoS appliance as the `admin` administrator, or an administrator account whose access profile has *Read-Write* permission for the *System* category.

Updating firmware on an HA pair requires additional steps. For details, see “Updating firmware on an HA pair” on page 87.

3. Go to *System* > *Status* > *Dashboard*.
4. In the *System Information* widget, in the *Firmware Version* row, click *Update*.

Figure 23:*System Information* widget

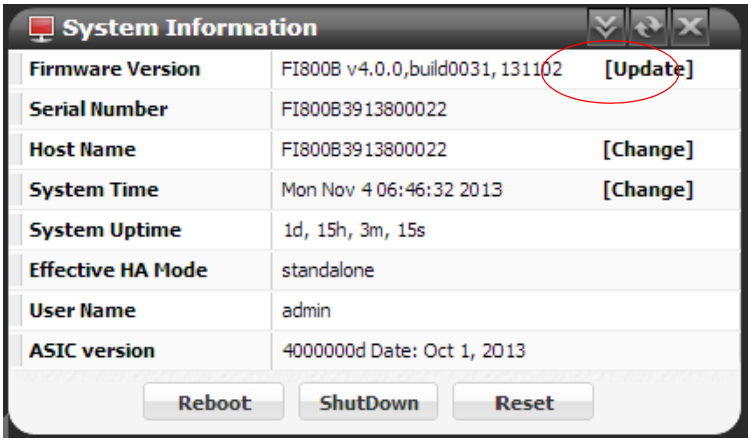

The *Firmware* dialog box is displayed.

5. Click *Browse* or *Choose File* (the button’s name varies by your browser) to locate and select the firmware file that you want to install, and then click *OK*.

6. Click *OK*.

   Your management computer uploads the firmware image to the FortiDDoS appliance. The FortiDDoS appliance installs the firmware and restarts. The time required varies by the size of the file and the speed of your network connection.

If you are ***downgrading*** the firmware to a previous version, and the settings are not fully backwards compatible, the FortiDDoS appliance may either remove incompatible settings, or use the feature's default values for that version of the firmware. You may need to reconfigure some settings.

When you upgrade or downgrade your firmware, the settings for the management network interface are reset to the defaults. To reconnect, either connect using a local connection and reconfigure MGMT 1 or use the default settings. See "Connecting to the web UI or CLI" on page 74.

7. Clear the cache of your web browser and restart it to ensure that it reloads the web UI and correctly displays all interface changes. For details, see your browser's documentation.
8. To verify that the firmware was successfully installed, log in to the web UI and go to *System* > *Status* > *Dashboard*.

   In the *System Information* widget, the *Firmware Version* row indicates the currently installed firmware version.
9. If you want to install alternate firmware on the secondary partition, follow "Installing alternate firmware" on page 87.
10. Continue with "Changing the "admin" account password" on page 93.

Installing new firmware overwrites any FortiGuard IP Reputation Service definitions. After any firmware update, re-enable the IP Reputation feature. FortiDDoS downloads the current definitions as part of the enabling process. For more information, see "FortiGuard IP Reputation Service" on page 130.

To install firmware via the CLI

1. Download the firmware file from the Fortinet Technical Support web site:

   https://support.fortinet.com/
2. Connect your management computer to the FortiDDoS console port using a RJ-45-to-DB-9 serial cable or a null-modem cable.

Updating firmware on an HA pair requires some additions to the usual steps for a standalone appliance. For details, see "Updating firmware on an HA pair" on page 87.

3. Initiate a connection from your management computer to the CLI of the FortiDDoS appliance, and log in as the `admin` administrator, or an administrator account whose access profile has *Read-Write* permission for the *System* category.

   For details, see "Connecting to the web UI or CLI" on page 74.
4. Connect MGMT 1 on the FortiDDoS appliance directly or to the same subnet as a TFTP server.
5. Copy the new firmware image file to the root directory of the TFTP server.

6. If necessary, start your TFTP server. (If you do not have one, you can temporarily install and run one such as `tftpd` (Windows, Mac OS X, or Linux) on your management computer.)

Because TFTP is ***not*** secure, and because it does not support authentication and could allow anyone to have read and write access, you should ***only*** run it on trusted administrator-only networks, ***never*** on computers directly connected to the Internet. If possible, immediately turn off `tftpd` off when you are done.

7. Verify that the TFTP server is currently running, and that the FortiDDoS appliance can reach the TFTP server.

   To use the FortiDDoS CLI to verify connectivity, enter the following command:

   ```
   execute ping 192.168.1.168
   ```

   where `192.168.1.168` is the IP address of the TFTP server.

8. Enter the following command to download the firmware image from the TFTP server to the FortiDDoS appliance:

   ```
   execute restore image tftp <filename_str> <tftp_ipv4>
   ```

   where `<filename_str>` is the name of the firmware image file and `<tftp_ipv4>` is the IP address of the TFTP server. For example, if the firmware image file name is `image.out` and the IP address of the TFTP server is 192.168.1.168, enter:

   ```
   execute restore image tftp image.out 192.168.1.168
   ```

   One of the following messages appears:

   ```
   This operation will replace the current firmware version!
   Do you want to continue? (y/n)
   ```

   or:

   ```
   Get image from tftp server OK.
   Check image OK.
   This operation will downgrade the current firmware version!
   Do you want to continue? (y/n)
   ```

9. Type `y`.

   The FortiDDoS appliance downloads the firmware image file from the TFTP server. The FortiDDoS appliance installs the firmware and restarts:

   ```
   MAC:00219B8F0D94
   ############################
   Total 28385179 bytes data downloaded.
   Verifying the integrity of the firmware image.
   Save as Default firmware/Backup firmware/Run image without
   saving:[D/B/R]?
   ```

If the download fails after the integrity check with the error message:

```
invalid compressed format (err=1)
```

but the firmware matches the integrity checksum on the Fortinet Technical Support web site, try a different TFTP server.

The time required varies by the size of the file and the speed of your network connection.

If you are ***downgrading*** the firmware to a previous version, the FortiDDoS appliance reverts the configuration to default values for that version of the firmware. You will need to reconfigure the FortiDDoS appliance or restore the configuration file from a backup. For details, see "Connecting to the web UI or CLI" on page 74 and, if you opt to restore the configuration, "Restoring a previous configuration" on page 169.

10. To verify that the firmware was successfully installed, log in to the CLI and type:

    ```
    get system status
    ```

    The firmware version number is displayed.
11. If you want to install alternate firmware on the secondary partition, follow "Installing alternate firmware" on page 87.
12. Continue with "Changing the "admin" account password" on page 93.

See also

- Updating firmware on an HA pair
- Installing alternate firmware

Updating firmware on an HA pair

Installing firmware on an HA pair is similar to installing firmware on a single, standalone appliance.

Start the update process with the master appliance, following the general firmware update instructions. Then, while the master appliance is rebooting, perform the update for the secondary appliance.

To ***downgrade*** to a previous version, switch out of HA, downgrade each appliance individually, and then switch them back into HA mode. If you attempt to downgrading while the appliances are still in HA mode, the HA daemon on the standby appliance might detect that the main appliance has older firmware and attempt to upgrade it to bring it into sync, undoing your downgrade.

See also

- Installing firmware
- Topology for synchronizing the configuration of two FortiDDoS appliances

Installing alternate firmware

You can install alternate firmware which can be loaded from its separate partition if the primary firmware fails. This can be accomplished via the web UI or CLI.

To install alternate firmware via the web UI

1. Download the firmware file from the Fortinet Technical Support web site:

   https://support.fortinet.com/

2. Log in to the web UI of the FortiDDoS appliance as the `admin` administrator, or an administrator account whose access profile has *Read-Write* permission for the *System* category.

Updating firmware on an HA pair requires slightly different steps than a standalone appliance. For details, see "Updating firmware on an HA pair" on page 87.

3. Go to *System* > *Maintenance* > *Backup & Restore*.

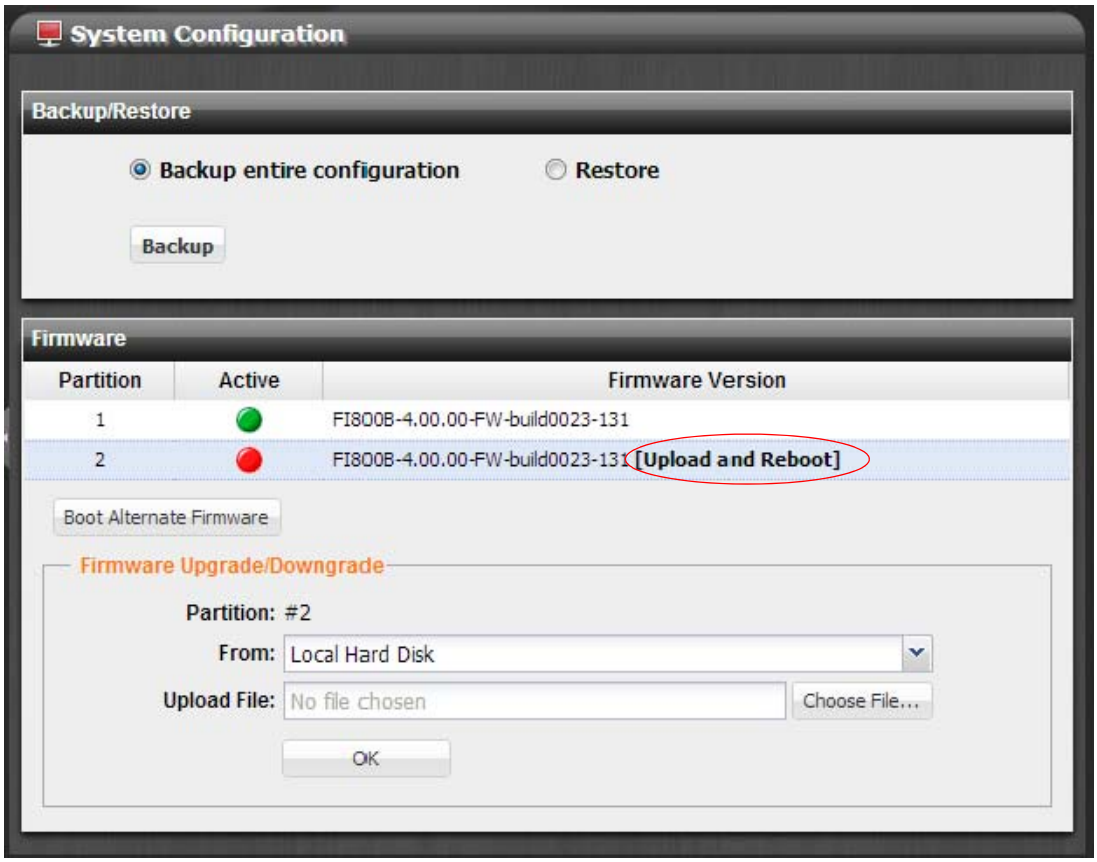

4. In the *Firmware* area, in the row of the alternate partition, click *Upload and Reboot*.

   The *Firmware Upgrade/Downgrade* dialog appears.
5. Click *Browse* to locate and select the firmware file that you want to install, and then click *OK*.
6. Click *OK*.

   Your management computer uploads the firmware image to the FortiDDoS appliance. The FortiDDoS appliance installs the firmware and restarts. The time required varies by the size of the file and the speed of your network connection.

If you are ***downgrading*** the firmware to a previous version, and the settings are not fully backwards compatible, the FortiDDoS appliance either removes incompatible settings or uses the feature's default values for that version of the firmware. You may need to reconfigure some settings.

7. Clear your web browser's cache, and then restart it to ensure that it reloads the web UI and correctly displays all interface changes. For details, see your browser's documentation.

8. To verify that the firmware was successfully installed, log in to the web UI and go to *System* > *Status* > *Dashboard*.

   In the *System Information* widget, the *Firmware Version* row indicates the currently installed firmware version.

To install alternate firmware via the CLI

1. Download the firmware file from the Fortinet Technical Support web site:

   https://support.fortinet.com/
2. Connect your management computer to the FortiDDoS console port using a RJ-45-to-DB-9 serial cable or a null-modem cable.
3. Initiate a connection from your management computer to the CLI of the FortiDDoS appliance, and log in as the `admin` administrator, or an administrator account whose access profile has *Read-Write* permission for the *System* category.

   For details, see "Connecting to the web UI or CLI" on page 74.
4. Connect MGMT 1 on the FortiDDoS appliance directly or to the same subnet as a TFTP server.
5. Copy the new firmware image file to the root directory of the TFTP server.
6. If necessary, start your TFTP server. (If you do not have one, you can temporarily install and run one such as `tftpd` (Windows, Mac OS X, or Linux) on your management computer.)

Because TFTP is ***not*** secure, and because it does not support authentication and could allow anyone to have read and write access, you should ***only*** run it on trusted administrator-only networks, ***never*** on computers directly connected to the Internet. If possible, immediately turn off `tftpd` off when you are done.

7. Verify that the TFTP server is currently running and that the FortiDDoS appliance can reach the TFTP server.

   To use the FortiDDoS CLI to verify connectivity, enter the following command:

   ```
   execute ping 192.168.1.168
   ```

   where `192.168.1.168` is the IP address of the TFTP server.
8. Enter the following command to restart the FortiDDoS appliance:

   ```
   execute reboot
   ```
9. As the FortiDDoS appliances starts, a series of system startup messages appear.

   ```
   Press any key to display configuration menu.....
   ```
10. Immediately press a key to interrupt the system startup.

You have only 3 seconds to press a key. If you do not press a key soon enough, the FortiDDoS appliance reboots and you must log in and repeat the `execute reboot` command.

If you successfully interrupt the startup process, the following messages appear:

```
[G]:  Get firmware image from TFTP server.
[F]:  Format boot device.
[B]:  Boot with backup firmware and set as default.
[Q]:  Quit menu and continue to boot with default firmware.
[H]:  Display this list of options.
```

```
Enter G,F,B,Q,or H:

Please connect TFTP server to Ethernet port "MGMT".
```

11. Type `G` to get the firmware image from the TFTP server.

    The following message appears:

    ```
    Enter TFTP server address [192.168.1.168]:
    ```

12. Type the IP address of the TFTP server and press Enter.

    The following message appears:

    ```
    Enter local address [192.168.1.188]:
    ```

13. Type a temporary IP address that can be used by the FortiDDoS appliance to connect to the TFTP server.

    The following message appears:

    ```
    Enter firmware image file name [image.out]:
    ```

14. Type the firmware image file name and press Enter.

    The FortiDDoS appliance downloads the firmware image file from the TFTP server and displays a message similar to the following:

    ```
    MAC:00219B8F0D94
    ############################
    Total 28385179 bytes data downloaded.
    Verifying the integrity of the firmware image.
    Save as Default firmware/Backup firmware/Run image without
    saving:[D/B/R]?
    ```

If the download fails after the integrity check with the error message:

```
invalid compressed format (err=1)
```

but the firmware matches the integrity checksum on the Fortinet Technical Support web site, try a different TFTP server.

15. Type `B`.

    The FortiDDoS appliance saves the backup firmware image and restarts. When the FortiDDoS appliance reboots, it is running the primary firmware.

See also

- Booting from the alternate partition
- Installing firmware

Booting from the alternate partition

System > *Maintenance* > *Backup & Restore* lists the firmware versions currently installed on your FortiDDoS appliance.

Each appliance can have up to two firmware versions installed. Each firmware version is stored in a separate partition. The partition whose firmware is currently running is noted with a white check mark in a green circle in the *Active* column.

To boot into alternate firmware via the web UI

1. Install firmware onto the alternate partition (see “Installing alternate firmware” on page 87).

2. Go to *System* > *Maintenance* > *Backup & Restore*.

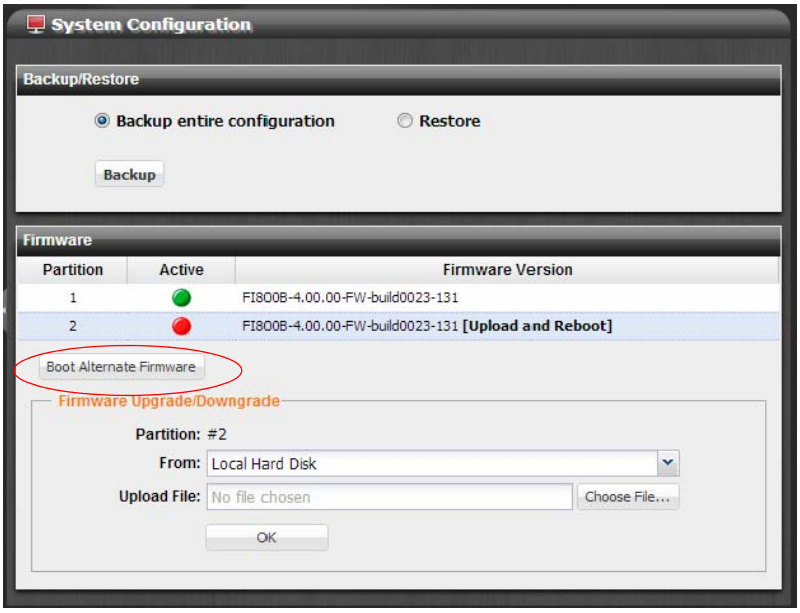

   To access this part of the web UI, ensure your administrator account's access profile has *Read-Write* permission for the *System* category. For details, see "Permissions" on page 44.

3. In the *Firmware* area, click *Boot Alternate firmware*.

   A warning message appears.

4. Click *OK*.

   A message appears instructing you to refresh your browser in a few minutes after the appliance has booted the other firmware.

To boot into alternate firmware via the local console CLI

1. Install firmware onto the alternate partition (see "Installing alternate firmware" on page 87).
2. Connect your management computer to the FortiDDoS console port using a RJ-45-to-DB-9 serial cable or a null-modem cable.
3. Initiate a connection from your management computer to the CLI of the FortiDDoS appliance, and log in as either the `admin` administrator or an administrator whose access profile has *Read-Write* permission for the *System* category.

   For details, see "Connecting to the web UI or CLI" on page 74.

4. Enter the following command to restart the FortiDDoS appliance:

   ```
   execute reboot
   ```

5. As the FortiDDoS appliances starts, a series of system startup messages appear.

   ```
   Press any key to display configuration menu........
   ```

   Immediately press a key to interrupt the system startup.

You have only 3 seconds to press a key. If you do not press a key soon enough, the FortiDDoS appliance reboots and you must log in and repeat the `execute reboot` command.

If you successfully interrupt the startup process, the following messages appear:

```
[G]:  Get firmware image from TFTP server.
[F]:  Format boot device.
[B]:  Boot with backup firmware and set as default.
[Q]:  Quit menu and continue to boot with default firmware.
[H]:  Display this list of options.

Enter G,F,B,Q,or H:

Please connect TFTP server to Ethernet port "MGMT".
```

6. Type `B` to reboot and use the backup firmware.

See also

- Installing alternate firmware

Changing the “admin” account password

The default administrator account, named `admin`, initially has no password.

Unlike other administrator accounts, the `admin` administrator account exists by default and cannot be deleted. The `admin` administrator account is similar to a root administrator account. This administrator account always has full permission to view and change all FortiDDoS configuration options, including viewing and changing all other administrator accounts. Its name and permissions cannot be changed.

Before you connect the FortiDDoS appliance to your overall network, you should configure the `admin` account with a password to prevent others from logging in to the FortiDDoS and changing its configuration.

Set a strong password for the `admin` administrator account, and change the password regularly. Failure to maintain the password of the `admin` administrator account could compromise the security of your FortiDDoS appliance. As such, it can constitute a violation of PCI DSS compliance and is against best practices.

To change the `admin` administrator password via the web UI

1. Go to *System* > *Admin* > *Administrators*.
2. In the row corresponding to the `admin` administrator account, double-click.

   The *Edit Administrator* dialog appears in a new panel below the list of accounts.
3. Mark the check box named *Change Password*.

   Additional text fields appear where you can enter the new password.

4. In the *Old Password* field, do not enter anything. (In its default state, there is no password for the `admin` account.)
5. In the *New Password* field, enter a password with sufficient complexity and number of characters to deter brute force and other attacks.
6. In the *Confirm Password* field, enter the new password again to confirm its spelling.

   The web UI alerts you with a red icon to the right of the field if the passwords do not match.
7. Click *Save*.

8. Click *Logout*.

   The new password takes effect the next time that administrator account logs in.

To change the `admin` administrator password via the CLI

Enter the following commands:

```
config system admin
  edit admin
    set password <new-password_str> ''
  end
exit
```

where `<new-password_str>` is the password for the administrator account named `admin`.

The new password takes effect the next time that the administrator account logs in.

Setting the system time & date

You can either manually set the FortiDDoS system time or configure the FortiDDoS appliance to automatically keep its system time correct by synchronizing with a Network Time Protocol (NTP) server. Using an NTP server is the recommended method.

The FortiDDoS system time must be accurate for many features to work, including scheduling, logging, traffic graphs, reports, and traffic statistics.

To prevent conflicts between traffic statistics recorded before and after the time settings change, it is recommended that you use *Protection Profiles* > *Factory Reset* > *Factory Reset* to reset traffic history for each profile.

To configure the system time via the web UI

1. Do one of the following:
   - Click *System* > *Maintenance* > *System Time*.

     The *Time Settings* dialog appears.

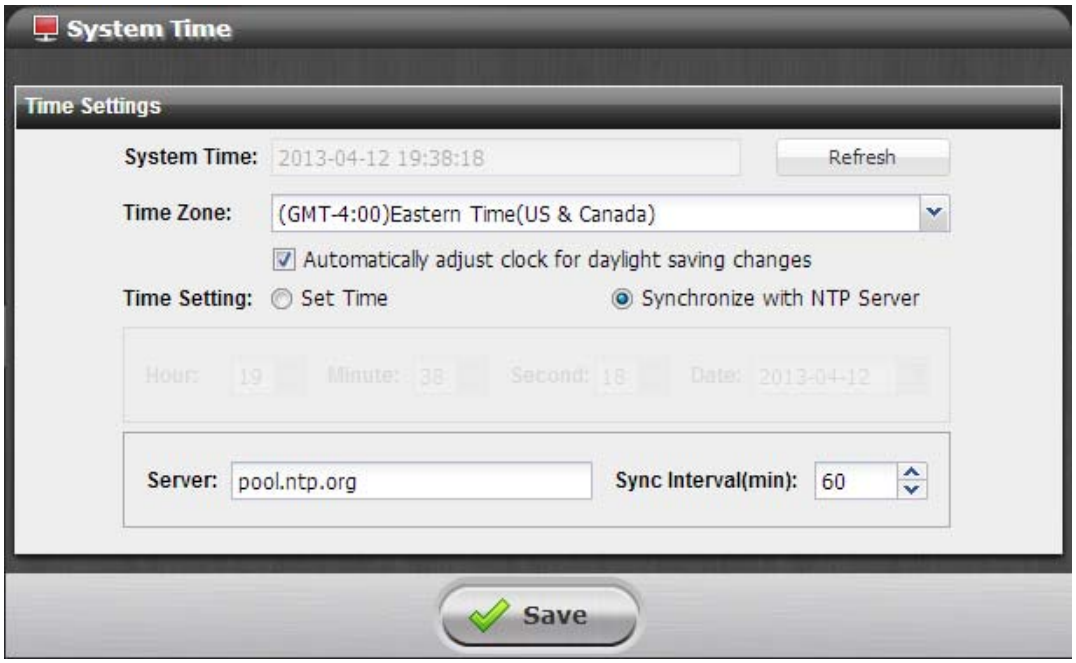

   - Go to *System* > *Status* > *Dashboard*. In the *System Information* widget, in the *System Time* row, click *Change*.

   To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *Maintenance* category. For details, see "Permissions" on page 44.
2. For *Time Zone*, select the time zone where the FortiDDoS appliance is located.

3. Do one of the following:
   - To configure FortiDDoS to automatically synchronize its clock with an NTP server, configure the following settings:

Setting name	Description
Synchronize with NTP Server	Select this option to automatically synchronize the date and time of the FortiDDoS appliance's clock with an NTP server, and then specify values for *Server* and *Sync Interval*.
Server	Type the IP address or domain name of an NTP server or pool, such as `pool.ntp.org`. To find an NTP server that you can use, go to http://www.ntp.org.
Sync Interval	Type how often the FortiDDoS appliance should synchronize its time with the NTP server, in minutes. For example, to synchronize its time once a day, enter 1440.

   - To manually set the current date and time, select *Set Time*. If you want FortiDDoS to automatically adjust its own clock when its time zone changes between daylight saving time (DST) and standard time, enable *Automatically adjust clock for daylight saving changes*. The clock is initialized with the time you specified when you click *Save*.
4. Click *Save*.

   You are prompted to confirm the change. The message also suggests that you reset the traffic history for each profile.

   If you manually configured the time, or if you enabled NTP and the NTP query for the current time ***succeeds***, the new clock time should appear in *System Time*. (If the query reply is slow, you may need to wait a couple of seconds, and then click *Refresh* to update the display in *System Time*.)

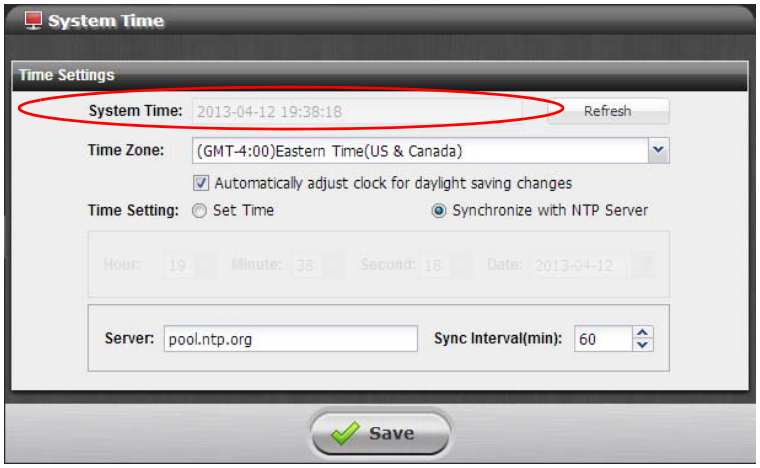

   If the NTP query ***fails***, the system clock continues without adjustment. For example, If FortiDDoS's time was 3 hours late, the time is still 3 hours late. Verify your DNS server IPs, your NTP server IP or name, routing, and that your firewalls or routers do not block or proxy UDP port 123.

To configure NTP via the CLI

To synchronize with an NTP server, enter the following commands:

```
config system time ntp
  set ntpsync enable
  set ntpserver {<server_fqdn> | <server_ipv4>}
  set syncinterval <minutes_int>
end
```

where:

- `{<server_fqdn> | <server_ipv4>}` is a choice of either the IP address or fully qualified domain name (FQDN) of the NTP server, such as `pool.ntp.org`
- `<minutes_int>` is how often in minutes the FortiDDoS appliance should synchronize its time with the NTP server

If your NTP query ***succeeds***, the new clock time should appear when you enter the command:

```
get system status
```

If the NTP query ***fails***, the system clock continues without adjustment. If FortiDDoS's time was 3 hours late, for example, the time is still 3 hours late. Verify your DNS server IPs, your NTP server IP or name, routing, and that your firewalls or routers do not block or proxy UDP port 123.

To manually set the date and time via the CLI

To manually configure the FortiDDoS appliance's system time and disable the connection to an NTP server, enter the following commands:

```
config system time ntp
  set ntpsync disable
end
config system time manual
  set zone <timezone_index>
  set daylight-saving-time {enable | disable}
end
```

where:

- `<timezone_index>` is the index number of the time zone in which the FortiDDoS appliance is located (to view the list of valid time zones and their associated index numbers, enter a question mark)
- `daylight-saving-time {enable | disable}` is a choice between enabling or disabling daylight saving time (DST) clock adjustments

See also

- System Information widget

Configuring network interfaces, gateway, and DNS

When shipped, the network interfaces used for management tasks have a default IP address and netmask. If these IP addresses and netmasks are not compatible with the design of your unique network, you must configure them.

Table 8: Default IP addresses and netmasks

Network Interface*	IP Address	Netmask
MGMT 1	192.168.1.99	255.255.255.0
MGMT 2	0.0.0.0	0.0.0.0

You also must configure FortiDDoS with the IP address of your DNS servers and gateway router.

Optionally, you can also select the link speed for the ports that receive incoming and outgoing traffic from network switches, routers or firewalls. By default, these ports use autonegociation. Change the speed only if the interface is connected to a device that does not support auto-negotiation.

If your appliance is protecting a network that has Internet Protocol version 6 (IPv6) traffic, enable IPv6 dual stack support (*Global Settings* > *Settings* > *Settings*).

You can use either the web UI or the CLI to configure these basic network settings.

See also

- Configuring the network interfaces
- Adding a gateway
- Configuring DNS settings
- Enabling Internet Protocol version 6 (IPv6) support

Configuring the network interfaces

You can configure network interfaces either via the web UI or the CLI.

If you plan to add this FortiDDoS to a FortiDDoS HA pair to synchronize their configuration, do ***not*** configure both mgmt1 and mgmt2. The HA heartbeat and synchronization link require a dedicated network interface. If you are re-cabling your network and must configure it, connect and switch to the new HA link ***first***. Failure to do so could cause unintentional downtime, failover, and ignored IP address configuration. To switch the HA link, see “Configuring configuration synchronization” on page 67.

To configure the IP address for management ports via the web UI

1. Go to *System* > *Network* > *Interface*.

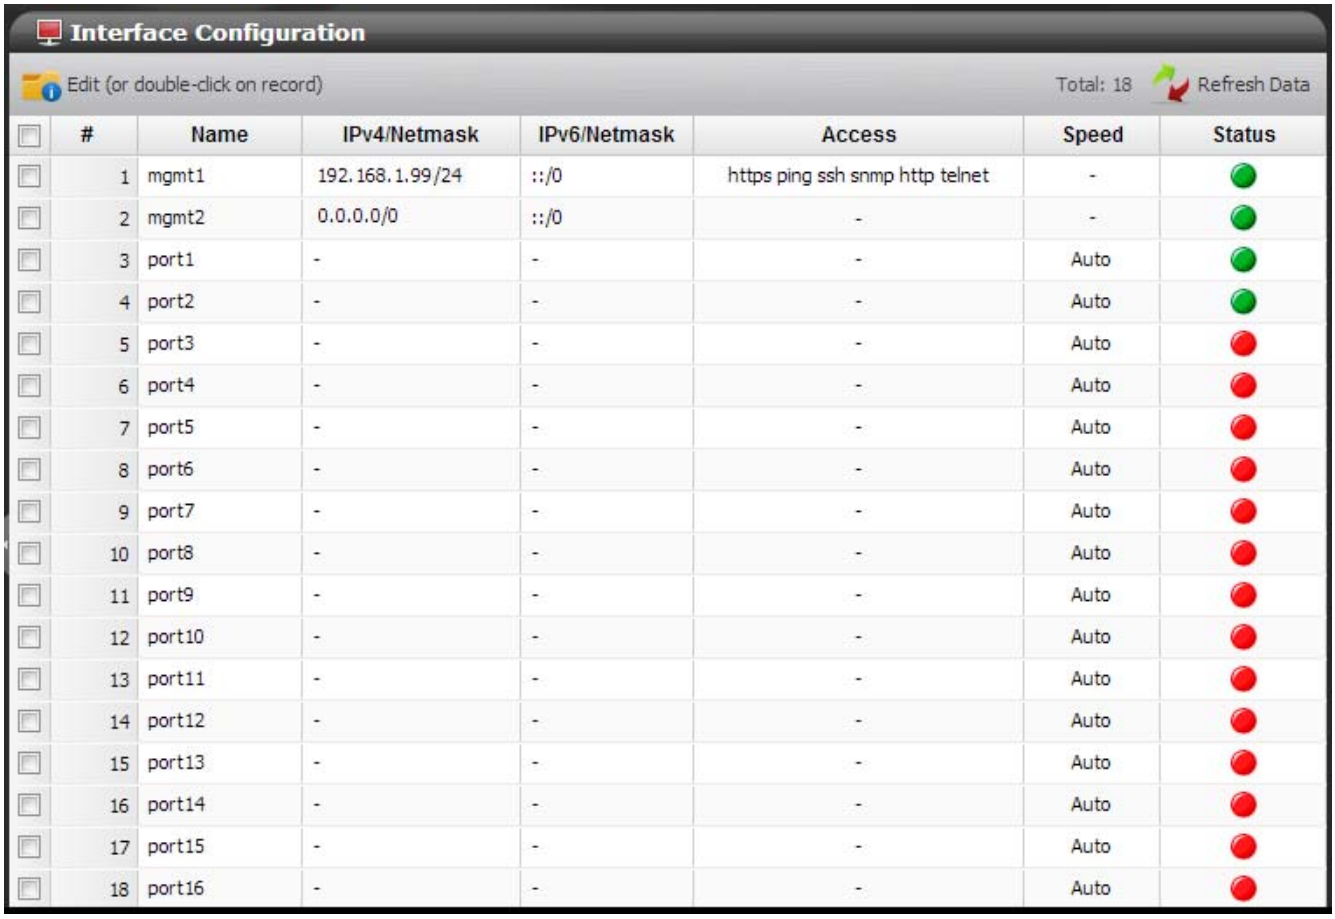

Interface Configuration

Edit (or double-click on record) Total: 18 Refresh Data

#	Name	IPv4/Netmask	IPv6/Netmask	Access	Speed	Status
1	mgmt1	192.168.1.99/24	::/0	https ping ssh snmp http telnet	-	
2	mgmt2	0.0.0.0/0	::/0	-	-	
3	port1	-	-	-	Auto	
4	port2	-	-	-	Auto	
5	port3	-	-	-	Auto	
6	port4	-	-	-	Auto	
7	port5	-	-	-	Auto	
8	port6	-	-	-	Auto	
9	port7	-	-	-	Auto	
10	port8	-	-	-	Auto	
11	port9	-	-	-	Auto	
12	port10	-	-	-	Auto	
13	port11	-	-	-	Auto	
14	port12	-	-	-	Auto	
15	port13	-	-	-	Auto	
16	port14	-	-	-	Auto	
17	port15	-	-	-	Auto	
18	port16	-	-	-	Auto	

To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *System* category. For details, see “Permissions” on page 44.

The Status indicators display the connectivity status. A green indicator means that the link is connected and negotiation was successful.

2. Double-click the *mgmt1* or *mgmt2* row.

   The dialog appears in a new panel below the list of interfaces.

3. Configure these settings:

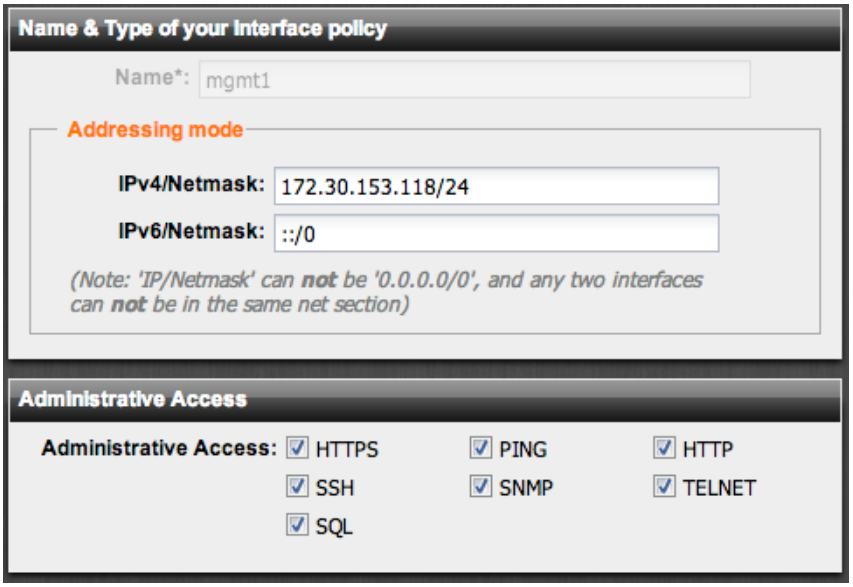

Setting name	Description
IPv4/Netmask	Type the IP address and CIDR-formatted subnet mask, separated by a forward slash (/), such as `192.0.2.5/24`. ***Dotted quad formatted subnet masks are not accepted.***
	The IP address must be on the same subnet as the network that the interface is connected to. Two network interfaces cannot have IP addresses on the same subnet.
IPv6/Netmask	Type the IP address and CIDR-formatted subnet mask, separated by a forward slash (/), such as `2001:0db8:85a3:::8a2e:0370:7334/64`. ***Dotted quad formatted subnet masks are not accepted.***
	The IP address must be on the same subnet as the network that the interface is connected to.
Administrative Access	Enable the types of administrative access that this interface permits.
	These options do ***not*** disable ***outgoing*** administrative connections, such as outgoing ICMP resulting from a CLI command such as `execute ping`. These options ***only*** govern ***incoming*** connections destined for the appliance itself.
	Caution: Enable ***only*** on network interfaces connected to trusted private networks (defined in *Trusted Host*) or directly to your management computer. If possible, enable only secure administrative access protocols such as HTTPS or SSH. Failure to restrict administrative access could compromise the security of your FortiDDoS appliance.
HTTPS	Enable to allow secure HTTPS connections to the web UI through this network interface. To configure the listening port number, see "Global web UI & CLI settings" on page 46.

Setting name	Description
PING	Enable to allow: • ICMP type 8 (`ECHO_REQUEST`)or type 30 for `traceroute` • UDP ports 33434 - 33534 for `ping` and `traceroute` to be received on this network interface. When it receives an `ECHO_REQUEST` (“ping”), FortiDDoS replies with ICMP type 0 (`ECHO_RESPONSE` or “pong”). **Note:** Disabling *PING* only prevents FortiDDoS from ***receiving*** ICMP type 8 (`ECHO_REQUEST`) or type 30 and traceroute-related UDP. It does ***not*** disable FortiDDoS CLI commands such as `execute ping` or `execute traceroute` that ***send*** such traffic.
HTTP	Enable to allow HTTP connections to the web UI through this network interface. To configure the listening port number, see “Global web UI & CLI settings” on page 46. **Caution:** HTTP connections are ***not*** secure, and can be intercepted by a third party. If possible, enable this option only for network interfaces connected to a trusted private network, or directly to your management computer. Failure to restrict administrative access through this protocol could compromise the security of your FortiDDoS appliance.
SSH	Enable to allow SSH connections to the CLI through this network interface.
SNMP	Enable to allow SNMP queries to this network interface, if queries have been configured and the sender is a configured SNMP manager. To configure the listening port number and configure queries and traps, see “SNMP traps & queries” on page 243. Disabling this setting does not disable outgoing SNMP traps.
TELNET	Enable to allow Telnet connections to the CLI through this network interface. **Caution:** Telnet connections are ***not*** secure, and can be intercepted by a third party. If possible, enable this option only for network interfaces connected to a trusted private network, or directly to your management computer. Failure to restrict administrative access through this protocol could compromise the security of your FortiDDoS appliance.
SQL	Enable to allow SQL access to the DDoS attack log through this network interface. For more information, see “Accessing the DDoS attack log using SQL” on page 231.

4. Click *Save*.

   If you were connected to the web UI through this network interface, you are now disconnected from it.

5. To access the web UI again, in your web browser, modify the URL to match the new IP address of the network interface. For example, if you configured the network interface with the IP address 10.10.10.5, you would browse to: `https://10.10.10.5`

   If the new IP address is on a different subnet than the previous IP address, and your computer is directly connected to the FortiDDoS appliance, you may also need to modify the IP address and subnet of your computer to match the FortiDDoS appliance’s new IP address.

To configure the IP address for a management port via the CLI

Enter the following commands:

```
config system interface
  edit {mgmt1|mgmt2}
    set ip <address_ipv4> <netmask_ipv4mask>
    set allowaccess {http https ping snmp ssh telnet}
  end
```

where:

- `{mgmt1|mgmt2}` is the interface to configure
- `<address_ipv4>` is the IP address assigned to the network interface
- `<netmask_ipv4mask>` is its netmask in dotted decimal format
- `{http https ping snmp ssh telnet}` is a space-delimited list of zero or more administrative protocols that you want to allow to access the FortiDDoS appliance through the network interface

HTTP and Telnet connections are ***not*** secure, and can be intercepted by a third party. If possible, enable this option only for network interfaces connected to a trusted private network, or directly to your management computer. Failure to restrict administrative access through this protocol could compromise the security of your FortiDDoS appliance.

If you were connected to the CLI through this network interface, you are now disconnected from it.

To access the CLI again, in your terminal client, modify the address to match the new IP address of the network interface. For example, if you configured the network interface with the IP address 10.10.10.5, you would connect to that IP address.

If the new IP address is on a different subnet than the previous IP address, and your computer is directly connected to the FortiDDoS appliance, you may also need to modify the IP address and subnet of your computer to match the FortiDDoS appliance's new IP address.

To configure the link speed using the web UI

By default, ports use autonegociation to determine the connection speed. Change the speed only if the interface is connected to a device that does not support autonegotiation.

If you are configuring interfaces for the FortiDDoS 1000B or FortiDDoS 2000B models, you can select *Auto* or *1000Mbps Full Duplex* only.

If the port connects to the network using a SFP transceiver (1 Gbps), set the speed to *1000Mbps Full Duplex*. Otherwise, set the speed to *Auto*.

1. Go to *System* > *Network* > *Interface*.

   To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *System* category. For details, see "Permissions" on page 44.
2. Double-click the row for the port you want to configure.
3. For *Speed*, select *Auto* or the appropriate speed and duplex.

To configure the link speed using the CLI

Enter the following commands:

```
config system interface
  edit <interface_name>
    set speed {auto | 10half | 10full | 100half | 100full | 1000half |
    1000full}
end
```

where:

- `<interface_name>` is the port to configure (`port1`, `port2`, `port3` and so on)
- `{auto | 10half | 10full | 100half | 100full | 1000half | 1000full}` is the speed and duplex for the port.

See also

- Configuring DNS settings
- Adding a gateway
- Global web UI & CLI settings

Adding a gateway

If your management computer is ***not*** directly attached to one of the physical ports of the FortiDDoS appliance, you may also require a static route that allows your management computer to connect with the web UI and CLI.

To add a static route via the web UI

1. Go to *System* > *Network* > *Static Route*.

   To access this part of the web UI, your administrator account's access profile must have *Read-Write* permission to items in the *Router* category. For details, see "Permissions" on page 44.
2. Click *Add*.

   A dialog appears.
3. Configure the following settings:

Setting name	Description
Interface	Select the network interface that uses the static route.
Destination IP/Mask	Type the destination IP address and network mask of packets that use this static route, separated by a slash (/) or space. The value `0.0.0.0/0` results in a default route, which matches all packets.
Gateway	Type the IP address of the next-hop router for the FortiDDoS management computer.

4. Click *Save*.

5. To verify that you can access FortiDDoS using the route, from a host on the route's destination network, attempt to connect to the FortiDDoS appliance's web UI via HTTP and/or HTTPS.

   If the connectivity test fails, you can use the CLI commands:

   ```
   execute ping <destination_ip4>
   ```

   to determine if a complete route exists from the FortiDDoS to the host, and

   ```
   execute traceroute <destination_ipv4>
   ```

   to determine the point of connectivity failure.

   Also enable *PING* on the FortiDDoS's network interface, or configure an IP address on the bridge, and then use the equivalent `tracert` or `traceroute` command on the host (depending on its operating system) to test the routing for traffic traveling in the opposite direction (from the host to the FortiDDoS).

   - If these tests fail, or if you do not want to enable *PING*, first examine the static route configuration on both the host and FortiDDoS.

     To display the routing table, enter the CLI command:

     ```
     diagnose netlink route list
     ```

     You may also need to verify that the physical cabling is reliable and not loose or broken, that there are no IP address or MAC address conflicts or blacklisting, and otherwise rule out problems at the physical, network, and transport layer.

   - If these tests ***succeed***, a route exists, but you cannot connect using HTTP or HTTPS, an application-layer problem is preventing connectivity.

     Verify that you have enabled *HTTPS* and/or *HTTP* on the management interface. Also examine routers and firewalls between the host and the FortiDDoS appliance to verify that they permit HTTP and/or HTTPS connectivity between them. Finally, you can also use the CLI command:

     ```
     diagnose system top 5 delay 30
     ```

     to verify that the daemons for the web UI and CLI, such as `sshd`, `cli`, `nginx`, and `php-fpm` are running and not overburdened.

To add a static route via the CLI

1. Enter the following commands:

   ```
   config system default-gateway
     edit <route_number>
       set destination <destination_ipv4/mask>
       set gateway <gateway_ipv4>
       set interface {mgmt1 | mgmt2}
     end
   ```

   where:

   - `<route_number>` is the number of the route in the list of static routes
   - `<destination_ipv4/mask>` specifies a dotted decimal IPv4 address and CIDR-notation netmask separated by a slash
   - `<gateway_ipv4>` is the IP address of the gateway router
   - `{mgmt1 | mgmt2}` specifies the network interface that uses the static route

   The FortiDDoS appliance should now be reachable to connections with networks indicated by the mask.

2. To verify connectivity, from a host on the route's destination network, attempt to connect to the FortiDDoS appliance's web UI via HTTP and/or HTTPS.

   If the connectivity test fails, you can use the CLI commands:

   ```
   execute ping
   ```

   to determine if a complete route exists from the FortiDDoS to the host, and

   ```
   execute traceroute
   ```

   to determine the point of connectivity failure. Also enable `ping` on the FortiDDoS (see "To configure the IP address for a management port via the CLI" on page 102), and then use the equivalent `tracert` or `traceroute` command on the host (depending on its operating system) to test the routing for traffic traveling in the opposite direction: from the host to the FortiDDoS.

   - If these tests ***fail***, or if you do not want to enable *PING*, first examine the static route configuration on both the host and FortiDDoS.

     To display all routes with their priorities, enter the CLI command:

     ```
     diagnose netlink route list
     ```

     You may also need to verify that the physical cabling is reliable and not loose or broken, that there are no IP address or MAC address conflicts or blacklisting, and otherwise rule out problems at the physical, network, and transport layer.

   - If these tests ***succeed***, a route exists, but you cannot connect using HTTP or HTTPS, an application-layer problem is preventing connectivity.

     Verify that you have enabled `http` and/or `http` on the management interface ("To configure the IP address for a management port via the CLI" on page 102). Also examine routers and firewalls between the host and the FortiDDoS appliance to verify that they permit HTTP and/or HTTPS connectivity between them. Finally, you can also use the CLI command:

     ```
     diagnose system delay 30
     ```

     to verify that the daemons for the web UI and CLI, such as `sshd`, `cli`, `nginx`, and `php-fpm` are running and not overburdened.

See also

- Configuring the network interfaces

Configuring DNS settings

Like many other types of network devices, FortiDDoS appliances require connectivity to DNS servers for DNS lookups.

Your Internet service provider (ISP) may supply IP addresses of DNS servers or you may want to use the IP addresses of your own DNS servers.

Incorrect DNS settings or unreliable DNS connectivity can cause issues with other features, including FortiGuard services and NTP system time.

For improved performance, use DNS servers on your local network.

You can use either the web UI or the CLI to configure the DNS settings. To verify your DNS settings, you use the CLI only.

To configure DNS settings via the web UI

1. Go to *System* > *Network* > *DNS*.

   To change settings in this part of the web UI, your administrator's account access profile must have *Write* permission to items in the *System* category. For details, see "Permissions" on page 44.
2. In *Primary DNS Server*, specify the IP address of the primary DNS server.
3. In *Secondary DNS Server*, specify the IP address of an alternative, secondary DNS server.
4. Click *Save*.

   The appliance queries the DNS servers whenever it needs to resolve a domain name into an IP address, such as for NTP system time or FortiGuard services.

   To verify the settings, see "To verify your DNS settings via the CLI" on page 106.

To configure DNS settings via the CLI

1. Enter the following commands:

```
config system dns
  set primary <address_ipv4>
  set secondary <address_ipv4>
end
```

   where:

   - `<address_ipv4>` is the IP address of a DNS server

   The appliance queries the DNS servers whenever it needs to resolve a domain name into an IP address, such as for NTP.

   To verify the settings, see "To verify your DNS settings via the CLI" on page 106.

To verify your DNS settings via the CLI

1. Enter the following commands:

```
execute traceroute <server_fqdn>
```

   where `<server_fqdn>` is a domain name such as www.example.com.

DNS tests may not succeed until you have completed "Adding a gateway" on page 103.

If the DNS query for the domain name ***succeeds***, you should see results that indicate that the host name resolved into an IP address, and the route from FortiDDoS to that IP address:

```
traceroute to www.example.com (192.0.43.10), 30 hops max, 46 byte
packets
 1  192.168.1.1 (192.168.1.1)  2.405 ms  0.629 ms  0.660 ms
 2  10.124.146.1 (10.124.146.1)  7.509 ms  8.822 ms  7.857 ms
 3  69.63.255.189 (69.63.255.189)  13.272 ms  18.270 ms  11.798 ms
 4  fallowfield2.cable.teksavvy.com (69.196.175.186)  8.872 ms
10.024 ms  8.624 ms
 5  fallowfield2.cable.teksavvy.com (69.196.175.185)  14.662 ms
13.030 ms  13.814 ms
...
11  43-10.any.icann.org (192.0.43.10)  32.059 ms  31.585 ms  32.127
ms
```

If the DNS query ***fails***, FortiDDoS displays an error message such as:

```
traceroute: bad address 'www.example.com'
```

Verify your DNS server IPs, routing, and that your firewalls or routers do not block or proxy UDP port 53.

See also

- Configuring the network interfaces
- Adding a gateway

Enabling Internet Protocol version 6 (IPv6) support

Enable *IPv6 dual stack support* if you plan to deploy FortiDDoS in a network that has Internet Protocol version 6 (IPv6) traffic.

When you enable IPv6 support, you must also specify the *IPv6 prefix length* and *IPv6 prefix* settings. FortiDDoS uses these settings to assign incoming IPv6 packets to the appropriate Service Protection Profile (SPP) and for destination tracking. For more information, see “Configuring the IPv6 prefix and prefix length settings” on page 108.

If you have service protection profile (SPP) policies that specify subnets using IPv4 addresses, delete and recreate these policies using IPv6 addresses after you enable and configure IPv6 dual stack support.

To enable IPv6 support via the web UI

1. Click *Global Settings* > *Settings* > *Settings*.
2. For *IPv6 dual stack support*, click *Enable*.
3. Complete the *IPv6 prefix* and *IPv6 destination tracking prefix* settings.

   For more information on these values, see “Configuring the IPv6 prefix and prefix length settings” on page 108.
4. Click *Save*.

To enable IPv6 support via the CLI

Enter the following commands:

```
config ddos global setting
  set ip-v6-dual-stack {enable | disable}
  set ip-v6-prefix-length {32 | 64 | 96}
  set ip-v6-prefix <ip_prefix>
end
```

where:

- `{enable | disable}` specifies whether IPv6 dual stack support is enabled
- `{32 | 64 | 96}` is the value of the prefix
- `<ip_prefix>` is the prefix of the network that FortiDDoS protects. Specify at least 64 bits of your network prefix.

See also

- Configuring the network interfaces
- Adding a gateway
- Configuring DNS settings

Configuring the IPv6 prefix and prefix length settings

FortiDDoS uses the IPv6 prefix and prefix length settings you specify for the following tasks:

- To assign incoming IPv6 packets to the appropriate Service Protection Profile (SPP)
- To provide a destination IP address for packets it generates and session diagnostics

The prefix values define 32-bit “windows” within the 128-bit value of the IPv6 address. When you select a prefix, you specify which portion of the address is fixed. FortiDDoS uses the 32-bit portion of the address that follows the fixed portion (that is, the “window”) to assign profiles.

The illustration shows these windows as blue, magenta and orange blocks.

Figure 24: IPv6 address “windows” for FortiDDoS configuration

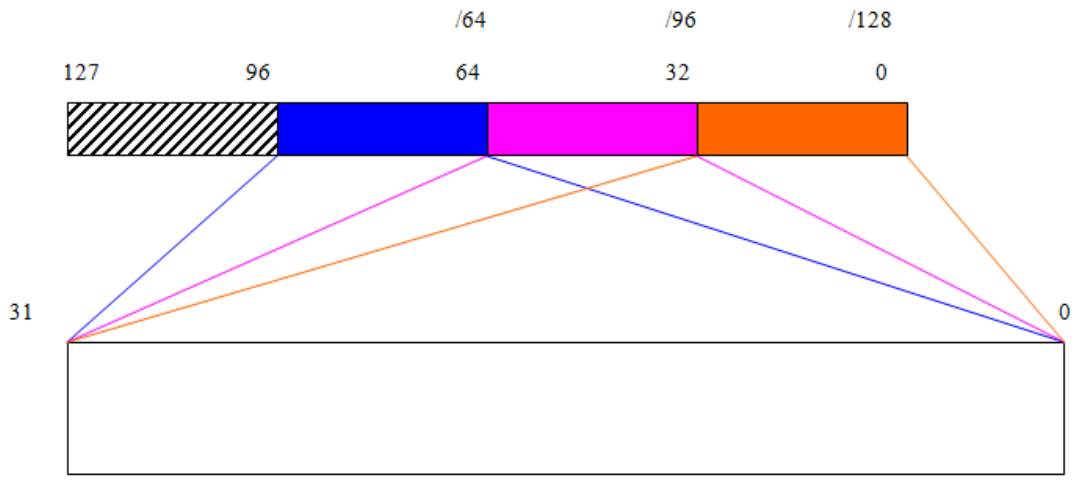

IPv6 prefix length

IPv6 prefix length identifies a 32-bit window within an IPv6 address that FortiDDoS uses to match IPv6 packet traffic to a Service Protection Profile (SPP).

For example, to assign servers and other resources to SPPs based on the portion of an address that is found in the 32-bit window that comes after the initial 64 bits of the address, for *IPv6 Prefix*, select */64*. Then, in the appropriate SPP policy configuration, for *IP address/Prefix*, define the subnet to protect by entering the complete address with a netmask value between 65 and 96. (For example, `9001:1234:5678:abcd:1122:3344:5566:7788/70`).

IPv6 prefix value	IP address bits used for SPP assignment
/32	33-64
/64	65-96
/96	97-128

Ensure that all SPP policy rules that specify subnets using IPv6 addresses use the same value for their initial 64 bits (for example, `9001:1234:5678:abcd`).

In addition, ensure that *IPv6 prefix* specifies the same initial 64 bits as the SPP policy rules. For example, `9001:1234:5678:abcd::/64`. (FortiDDoS does not validate the values that you specify.)

IPv6 prefix

FortiDDoS uses the *IPv6 prefix* value to generate destination IP address information for the following features:

- Packets generated by the appliance. For example, SYN and ACK cookies used for SYN flood mitigation, or RST and ACK packets used for aggressive aging.
- Session diagnostics.

You can specify more than 64 bits of the IP address, but FortiDDoS ignores any values beyond the initial 64 bits.

Ensure that the initial 64 bits specified in *IPv6 prefix* are the same as those specified by *IP address/Prefix* values in any SPP policy rules that specify subnets using IPv6 addresses. (For example, if *IPv6 prefix* is `9001:1234:5678:abcd::/64`, a valid *IP address/Mask* value is `9001:1234:5678:abcd:1122:3344:5566:7788/70`.)

Identifying IP addresses and subnets to protect (SPP creation)

The Service Protection Profile (SPP) functionality enables FortiDDoS to behave as if it were multiple physical appliances, with each appliance protecting a single server or group of networked servers. You can build an SPP that is customized for the characteristics of each system you want to protect.

For general information on the features and advantages of SPPs, see "Service Protection Profiles (SPPs)" on page 41.

FortiDDoS maintains traffic history for each Service Protection Profile (SPP). It uses this data to generate recommended thresholds, dynamically adjust thresholds, and generate traffic statistics.

Any change to a profile's configuration or the resources it monitors has an effect on this traffic history and, in turn, the threshold estimation mechanism and traffic graphs. For example, removing a server from the profile, changing which servers are in the profile, or changing the services offered by a monitored server.

It is strongly recommended that you reset the traffic history for a profile before you make any significant changes to its configuration. To reset the profile, go to *Protection Profiles* > *Factory Reset* > *Factory Reset*.

SPP Policy configuration

You use the SPP *Config* settings to create a profile by specifying its name and ID.

The *SPP Policy* settings are a set of one or more rules. Each SPP rule assigns a subnet (specified by an IP address and netmask and subnet ID) to the Service Protection Profile that protects it. One rule is reserved for the default profile (SPP-0) and you can add up to 511 additional rules to each FortiDDoS appliance.

All traffic that has a source or destination IP address that matches the rule's IP address and netmask is monitored and regulated by the associated Service Protection Profile.

The rule name helps you to remember information about the subnet or its configuration, such as the name of the server that the profile protects.

The subnet ID identifies information that is specific to the subnet for reports. (To view an automatically generated report of the top attacked subnets, go to *Log & Report* > *Report Browse* > *Executive Summary*. For information on generating a report that contains attack information by subnet, see "Configuring a report" on page 250.)

You can associate multiple rules with each Service Protection Profile. For example, you can create five rules, each specifying a subnet that represents a network, but configure the rules to use only two profiles: three different networks assigned to SPP-1 and the two other networks assigned to SPP-2.

SPP Policy rule priority

FortiDDoS matches traffic to rules by searching the SPP Policy entries top-to-bottom. If a packet matches more than one rule, FortiDDoS uses the rule closest to the top of the list.

Because you can specify a subnet definition in a SPP Policy rule as a range, ensure that any rules that apply to specific networks within the range are higher in the list than the ranges that they are part of.

Although FortiDDoS adds items to the list in the order that you create them, you can change the order of items by dragging.

Default SPP

By default, FortiDDoS has a service protection profile (SPP) named SPP-0, which you cannot delete, rename or assign subnets to.

FortiDDoS uses SPP-0 to monitor and regulate the following types of packets:

- Packets that do not match a subnet in *Global Settings* > *Service Protection Profiles* > *SPP Policy*.
- Packets that have a corrupted IP header.

Because FortiDDoS generates reports by SPP and subnet, you can access information about these packets by generating or viewing reports for SPP-0.

SPP-0 is a catch-all profile and its traffic statistics are affected by the traffic that FortiDDoS assigns to is by default. Therefore, to maintain accurate traffic statistics for your subnets, do not associate any subnets with SPP-0.

SPP and subnet names and IDs

When you create an SPP, the names you specify for the profiles and subnets help you identify them in other parts of the UI, including attack event log messages.

Specify subnets to make it easier to identify which network resources are under attack. If you do not specify the subnets within an SPP, FortiDDoS identifies the attacked entity by profile only and not the individual subnet.

For example, create an SPP named `web_servers`, one named `DNS_servers`, and so on.

For the *SPP Policy* rules that define subnets for each profile, you can assign a name such as `Web_Server_Yosemite` to a specific IP address. (FortiDDoS uses the subnet ID assigned in the rule to identify the IP address in the reports that provide attack statistics by subnet.)

Switching SPPs automatically

You can specify an alternate Service Protection Profile for each subnet. FortiDDoS automatically switches to the alternate profile when the level of traffic exceeds a threshold that you specify.

For example, you can automatically switch from a profile that is designed to handle low levels of traffic to another, more stringent profile when traffic exceeds a specified threshold. Or you can switch from an SPP running in detection mode to one that uses prevention mode when traffic crosses a threshold.

FortiDDoS monitors and regulates traffic for the subnet using the alternate profile as long as the traffic level exceeds the threshold. After traffic is less than the threshold and remains below it for a timeout period that you specify, FortiDDoS switches back to using the primary profile configuration to protect the subnet.

You enable SPP Switching Policy and specify a timeout period for all profiles before you enable the feature for an individual profile.

For information on enabling the feature for an individual profile, see “Create a service protection profile (SPP)” on page 113.

To configure FortiDDoS to generate a log message whenever it switches a subnet to its alternate SPP, see “Selecting which system events to log” on page 236. If you want FortiDDoS to simply notify you that the traffic level has exceeded the SPP switching threshold without switching the SPP, in the SPP policy settings, specify the same SPP for both *Service Protection Profile* and *Alternate Service Protection Profile*.

To configure FortiDDoS to send an email message to the configured recipient whenever it switches a subnet to its alternate SPP, see “Alert email” on page 241.

To configure the SPP switching feature via the web UI

1. Enable the feature for the appliance and specify the timeout used by all SPPs using *Global Settings* > *Service Protection Profiles* > *Switching Policy*.

   The timeout value is specified in seconds. The default and maximum value is 255.
2. Configure the feature for individual SPPs using *Global Settings* > *Service Protection Profiles* > *SPP Policy*.

To configure the SPP switching feature via the CLI

Enter the following commands:

```
config ddos global spp-switching-policy
  set Switching {enable | disable}
  set timeout <integer>
end
```

where:

- `{enable | disable}` specifies whether the feature is enabled
- `<integer>` specifies a timeout in seconds. When the alternate profile is in use and traffic volume is lower than the switching threshold for this period of time, FortiDDoS returns to using the normal profile to monitor and protect the subnet.

See also

- Create a service protection profile (SPP)
- Selecting which system events to log
- Alert email

Create a service protection profile (SPP)

To create a service protection profile (SPP) using the web UI

1. Click *Global Settings* > *Service Protection Profiles* > *Config.*

   To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *Global Settings* category. For details, see “Permissions” on page 44.
2. Click *Add* and enter the following information:

Setting name	Description
Name	The name of the Service Protection Profile. This name is used in other parts of the FortiDDoS configuration. A profile’s logical name can help you remember its characteristics. For example, you can name one profile `web_servers`, another profile `DNS servers`, and so on. The name cannot contain spaces.
ID	Select a number between 1 and 7 that identifies this profile. (0 is reserved for SPP-0, the default profile)

3. Click *Save*.
4. Click *Global Settings* > *Service Protection Profiles* > *SPP Policy.*

5. Click *Add*, and then configure the following settings:

Setting name	Description
Name	Enter a name that describes the subnet.
Subnet ID	Specify a value between 1 and 511 that identifies the subnet. The subnet ID identifies information that is specific to the subnet for reports. (See "SPP and subnet names and IDs" on page 112.)
IP address/Mask/ IP address/Prefix	Specify the IP address of the subnet. All traffic with a source or destination IP address that matches this address is protected by the specified Service Protection Profile. Type the IP address and CIDR-formatted subnet mask, separated by a forward slash (/). If *IPv6 dual stack support* (*Global Settings* > *Settings* > *Settings*) is enabled, specify an Internet Protocol version 6 (IPv6) address and prefix. Ensure that the address has the same intial 64 bits that are specified by *IPv6 prefix* (*Global Settings* > *Settings* > *Settings*). For more information, see "Configuring the IPv6 prefix and prefix length settings" on page 108. Ensure that all SPP policy rules that specify subnets using IPv6 addresses use the same value for their initial 64 bits (for example, `9001:1234:5678:abcd`).
Service Protection Profile	Select the profile that protects the specified subnet. For example, the profile name you specified earlier. SPP-0 is a catch-all profile and its traffic statistics are affected by the traffic that FortiDDoS assigns to is by default. Therefore, to maintain accurate traffic statistics for your subnets, do not associate any subnets with SPP-0. For more information, see "Default SPP" on page 112.
Enable SPP Switching	When this option is enabled and subnet traffic exceeds the value specified by *Threshold*, the profile specified by *Alternate Service Protection Profile* protects the subnet. To enable this feature and specify the timeout value it uses, click *Global Settings* > *Service Protection Profiles* > *Switching Policy*. To receive notification when FortiDDoS switches profiles, enable system event logging or alert email for SPP switching.
Alternate Service Protection Profile	(Available only when SPP switching is enabled.) Specifies the profile that protects the subnet when its traffic exceeds the value specified by *Threshold*. If you want FortiDDoS to simply notify you that the traffic level has exceeded the SPP switching threshold without switching the SPP, specify the same SPP value as *Service Protection Profile*.
Threshold	(Available only when SPP switching is enabled.) Specifies the maximum packet rate for the specified profile. When traffic exceeds this rate, the profile specified by *Alternate Service Protection Profile* protects the subnet.
Comments	An additional description or notes related to the rule, as required.

6. Click *Save*.

7. To assign additional subnets to the profile, add additional *SPP Policy* rules.

   To change the order of the rules in the list, click and hold an item, and then drag it to its new position.

For more information on configuring profiles, see "Service Protection Profile settings" on page 177.

To create a service protection profile (SPP) using the CLI

For additional information on the settings, see "To create a service protection profile (SPP) using the web UI" on page 113."

1. Enter the following commands:

```
config spp
   edit <spp_name>
end
```

   where `<spp_name>` is the name of the profile to configure.

2. Enter the following commands:

```
config ddos global spp-policy
   edit <rule_name>
      set subnet-id <entry_index>
      set ip <address_ip/mask>
      set spp <spp_name>
end
```

   where:

   - `<rule_name>` specifies a name that identifies the rule in the list
   - `<entry_index>` is the index number of the item in the list of rules
   - `<address_ip/mask>` is an IP address and CIDR-formatted subnet mask, separated by a forward slash (/)
   - `<spp_name>` specifies the profile that monitors and regulates the subnet specified by the IP address and subnet mask

3. To change the order of rules in the policy list, enter the following commands:

```
config ddos global spp-policy
   move <entry_index> after <entry_index>
end
```

   where `<entry_index>` is the index number of the items in the list of rules.

See also

- Creating an SPP for UDP traffic
- Identifying IP addresses and subnets to protect (SPP creation)
- Switching SPPs automatically
- Selecting which system events to log
- Alert email

Creating an SPP for UDP traffic

The only mitigation mechanisms FortiDDoS currently provides for attacks via UDP traffic are source tracking and rate limiting. The source tracking feature allows FortiDDoS to control an attack from a single source or a limited number of sources. However, if the attack is distributed,

the rate limiting feature limits all UDP traffic in the service protection profile (SPP), including legitimate traffic. It cannot limit only the UDP traffic that is associated with the attack.

Unlike TCP traffic, FortiDDoS currently does not track the state of UDP and DNS traffic. This means that for UDP traffic, it does not differentiate requests from responses. If you set thresholds for UDP ports that are too low - either manually or by using the system recommended values - FortiDDoS may block harmless UDP traffic and generate drops that are false positives.

For example, in a DNS request, the destination UDP port is 53 and the source port is a randomly chosen UDP port. The response uses a source port of 53 and the destination port is the source port from the original request, which creates a lot of outbound traffic on many destination UDP ports. If you lower the traffic thresholds for ports other than 53, and the operating mode for outgoing traffic is *Prevention*, FortiDDoS may block and rate limit responses.

To avoid these types of false positives, create a Service Protection Profile (SPP) that regulates UDP traffic only and set its UDP port thresholds to a reasonable value. If the operating mode for outgoing traffic is *Prevention*, set all outbound thresholds for UDP ports to a high or reasonably high value.

Alternatively, to ensure continued service, you can change your DNS service to a third-party provider.

For information on creating an SPP, see “Identifying IP addresses and subnets to protect (SPP creation)” on page 111.

For information on specifying the values for UDP port thresholds, see “Adjusting thresholds individually” on page 155 and “Specifying Protocols, TCP Ports, and UDP Ports thresholds” on page 166.

Setting FortiDDoS to detection mode

For each Service Protection Profile (SPP), FortiDDoS operates in either detection or prevention mode. In most cases, you first configure FortiDDoS to operate in detection mode and then switch it to prevention mode after an initial learning period.

You use *Protection Profiles* > *SPP Settings* > *SPP Settings* to specify an operating mode for inbound and outbound traffic for each Service Protection Profile (SPP).

Detection mode

In detection mode, FortiDDoS logs events and builds traffic statistics for the profile but does not limit or block any data packet traffic. Packets pass through the FortiDDoS as they travel to and from one or more protected systems and the rest of the network.

In most cases, you use this mode for an initial learning period of 2-14 days.

Ensure that there are no attacks during the initial learning period and that it is long enough to be a representative period of activity. If activity is heavier in one part of the week than another, ensure that your initial learning period includes periods of both high and low activity. In most cases, weekends alone do not make a good baseline for weekday activity.

In most cases, a seven-day learning period provides a reasonable profile of “normal” traffic with weekly variations. After you switch it to prevention mode, FortiDDoS continues to learn traffic patterns and continuously adjusts traffic profiles.

Prevention mode

Prevention mode is FortiDDoS's fully functional operating mode. When a profile is in prevention mode, the appliance limits or blocks any anomalous traffic or traffic that exceeds threshold values on the subnets associated with the profile.

To set the operating mode using the web UI

1. Click *Protection Profiles* > *SPP Settings* > *SPP Settings*.
2. In the top-right corner of the content pane, for *Service Protection Profiles*, select the profile to configure.
3. For *Inbound operating mode*, select *Detection* or *Prevention*.
4. For *Outbound operating mode*, select *Detection* or *Prevention*.
5. Click *Save*.

To set the operating mode via the CLI

Enter the following commands:

```
edit <spp_name>
   config ddos spp setting
      set inbound-operating-mode {detection | prevention}
      set outbound-operating-mode {detection | prevention}
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `{detection | prevention}` specifies the operating mode to use

See also

- Configuration workflow
- Setting thresholds to system recommended values
- Avoiding disruptions while adjusting thresholds
- System Status widget

Customizing protection features for protected subnets

The protection features and setting that you select for a Service Protection Profile (SPP) depend on the characteristics of the server or other network resource that the profile protects. Servers can handle many different types of traffic: SIP, SSL, port 80, and so on. In addition, networks can be prone to specific attacks such as SYN floods and botnet attacks. Before any learning begins, you should customize thresholds and ACLs for the IP addresses and address ranges you have specified.

For example, if a server does not handle SIP traffic, block this type of traffic using the ACL. If your network includes DNS servers, you can use a dedicated logical protection zone (Service Protection Profile or SPP) to protect them.

See also

- Access control lists (ACLs)
- FortiGuard IP Reputation Service
- Enabling higher thresholds for proxy server IP addresses
- Do Not Track Policy list
- Configuring SYN flood mitigation feature controls
- SYN flood and zombie flood prevention
- Configuring SYN flood mitigation feature controls

Preset access control

FortiDDoS can automatically drop packets with the following characteristics, which usually indicate attack traffic. These settings are preset: you can enable or disable them but cannot configure them. FortiDDoS does not report the packets that it drops because of these characteristics in traffic graphs or reports:

- IP first fragments for IPv4 and IPv6 packets that do not pass a Dos attack check
- TCP packets with control flags = 0 and sequence number = 0
- TCP SYN packets with source port between 0-1023
- Packets with MAC SA = DA
- IPv4 and IPv6 packets where the SIP = DIP
- IPv6 fragments smaller than the minimum size
- Fragmented ICMP packets
- TCP fragments with offset value of 1
- UDP and TCP packets with Sport = Dport
- TCP packets with SYN and FIN bits set
- TCP packets with FIN, URG and PSH bits set and Seq number = 0

You enable preset access control using the command line interface (CLI).

To turn the built-in access control on or off

Execute the following commands:

```
execute dos-control {enable | disable}
```

Access control lists (ACLs)

FortiDDoS provides access control lists (ACLs) that deny or allow packet traffic based on criteria such as IP address, port number, protocols, URLs or HTTP header field values.

FortiDDoS provides the following types of ACLs:

- A global ACL that controls packet traffic for all Service Protection Profiles (SPPs), based on the following criteria:
  - IP-Netmask (IPv4)
  - IP-Address (IPv4)
  - Geolocation
  - IP-Netmask (IPv6)
  - IP-Range (IPv6)
- SPP ACLs that control packet traffic for individual profiles, based on the following criteria:
  - IP Address
  - Fragment
  - Protocol
  - TCP Port
  - UDP Port
  - ICMP Type/Code
  - URL
  - HTTP header field: Host, Referer, Cookie, User Agent
- A global ACL that blocks access to all Service Protection Profiles (SPPs) for specified protocols.

Blocking dark and bogon addresses

There are many IP addresses that should not appear as either source or destination addresses in an IP packet.

A bogon is an informal name for an IP packet on the public Internet that claims to be from an area of the IP address space that is reserved but not yet allocated or delegated by the Internet Assigned Numbers Authority (IANA) or a delegated Internet registry. The areas of unallocated address space are called “bogon space” or “dark address space”.

The term “bogon” stems from hacker jargon, where it is defined as the quantum of “bogosity”, or the property of being bogus. A bogon packet is frequently bogus both in the conventional sense of being forged for illegitimate purposes, and in the hackish sense of being incorrect, absurd, and useless.

In a private network, this could mean undefined private addresses should not be expected as source or destination. For example, if an enterprise uses only the 192.168.3.x range within its private domain, and then any other private addresses such as 192.168.1.x, 192.168.2.x and 192.168.4.x through 92.168.254.x are illegal. Use of these addresses usually means stealth activity that is mostly performed by worms.

In a public network, this would mean all bogon-prefixes should not appear as source or destination. A bogon prefix is a route that should never appear in the Internet routing table. A packet routed over the public Internet (not including over VPN or other tunnels) should never have a source address in a bogon range. These are commonly found as the source addresses of DDoS attacks.

Bogon prevention is a component of anti-spoofing.

The following site is a good source of bogon addresses:

http://www.cymru.com/Documents/bogon-dd.html

You can configure FortiDDoS to block these types of addresses by adding them to its global ACL or the ACL for an individual profile.

Examples:

- To deny spoofing packets from the Internet with the source address 192.168.x.x, add IP Address 192.168.0.0, Netmask 255.255.0.0.
- To deny outbound spoofing packets (that is, to deny addresses that are not in your inside LAN) with the source address as private addresses 172.16.x.x, add IP Address 172.16.0.0, Netmask 255.255.0.0.
- To deny the address range 10.x.x.x altogether because it is "dark" both inside and outside your network, add IP Address 10.0.0.0, Netmask 255.0.0.0.

For instructions on how to block addresses, see "Add addresses or locations to the global ACL" on page 122 and "Access and tracking control for Service Protection Profiles" on page 124.

Specify addresses that can exceed thresholds (whitelist)

In addition to denying harmful IP addresses, you can specify IP addresses that are known to be good and are always allowed, even if they exceed set thresholds. This group of addresses is sometimes called a whitelist.

For example, you can add addresses for devices that perform backups, which can have a high traffic profile because they need to establish many connections or send a large number of packets to perform their tasks.

Packets from IP addresses that are denied or allowed by ACLs do not affect the statistics for continuous learning for source addresses. However, other characteristics of the packets, such as protocols and ports, affect the corresponding statistics.

You cannot add IP-Netmask (IPv4) or Geolocation addresses as whitelist items in the global ACL.

FortiDDoS does not track any connections for items that are configured in the global ACL as *Allow*. However, it does track the source and associated traffic for items that are configured as *Track & Allow* in the ACL for an SPP.

For instructions on how to specify addresses that are always allowed, see "Add addresses or locations to the global ACL" on page 122 and "Access and tracking control for Service Protection Profiles" on page 124.

Blocking addresses from a specific geographic location, anonymous proxies, and satellite providers

FortiDDoS allows you to block traffic from geographic locations. For example, you can block traffic from locations where you do not have any customers or that are unlikely to generate traffic. This type of filtering can significantly reduce unnecessary traffic to servers, even during periods of normal activity.

The Geolocation address option also allows you to block anonymous proxy and satellite provider addresses.

You can configure blocking by geolocation for IP version 4 (IPv4) addresses and global ACLs only. The *Geo Location* option is not available for IP version 6 (IPv6) address configuration.

FortiDDoS captures information about packets from denied locations in the following locations:

- Traffic graphs: *Monitor* > *Aggregate ACL Drops* > *Layer 3*, *Monitor* > *Layer 3* > *Denied Countries*
- Executive Summary dashboard (*Log & Report* > *Report Browse* > *Executive Summary*)
- Subnet Executive Summary dashboard (*Log & Report* > *Report Browse* > *Subnet Executive Summary*)
- Reports

To block addresses from a geographic location

1. Go to *Global Settings* > *Settings* > *Settings* and do one of the following:
   - To permit traffic from all geographic locations and from anonymous proxies and satellite providers, with a few exceptions you specify, select *Allow all and deny some*.
   - To block traffic from all geographic locations and from anonymous proxies and satellite providers, with a few exceptions you specify, select *Deny all and allow some.*
2. To configure the exceptions to the *Geo Location Policy* setting, go to *Global Settings* > *Address* > *Address Config*.
3. To add the exceptions to the global ACL, go to *Global Settings* > *Access Control List* > *Access Control List.*

For more information on configuring exceptions, see “Add addresses or locations to the global ACL” on page 122.

Blocking specific protocols

The ACL for each service protection profile (SPP) allows you to block all traffic associated with a specific protocol. However, in some cases, it is more efficient to block specific protocols before FortiDDoS evaluates the packets using the SPP configuration.

For example, to prevent large-scale, brute force attacks using UDP or ICMP, you can configure FortiDDoS to block these protocols on a subnet that is monitored by more than one SPP.

FortiDDoS does not report the packets that it drops because of settings in this ACL in traffic graphs or reports.

You configure this global ACL for protocols using the CLI only. For instructions, see “Blocking a protocol for a specified subnet” on page 128.

See also

- Add addresses or locations to the global ACL
- Access and tracking control for Service Protection Profiles
- Creating an IP address item to use with ACLs
- Creating a service item to use with ACLs
- Add an address or service to a profile's ACL
- Blocking a protocol for a specified subnet

Add addresses or locations to the global ACL

When you add an address to the global ACL that is type *Geo Location*, you are specifying an exception to the value of *Geo Location Policy* setting. For more information, see "Blocking addresses from a specific geographic location, anonymous proxies, and satellite providers" on page 120.

To add addresses or locations to the global ACL using the web UI

1. Do one of the following:
   - To configure Internet Protocol version 4 (IPv4) addresses, click *Global Settings* > *Address* > *Address Config*.
   - To configure Internet Protocol version 6 (IPv6) addresses, click *Global Settings* > *Address* > *Address Config (IPv6)*.
2. Click *Add*.
3. Enter a name that identifies the address in the ACL.
4. Select the type:

IP-Netmask	Specifies a subnet using an Internet Protocol version 4 address and netmask.
IP-Address	Specifies a subnet using an Internet Protocol version 4 address.
Geo Location	Allows you to select a geographic location, *Anonymous Proxy* or *Satellite Provider*. Available for IPv4 addresses only.
IPv6-Netmask	Specifies a subnet using an Internet Protocol version 6 address and netmask.
IPv6-Address	Specifies a subnet using an Internet Protocol version 6 address.

5. Do one of the following:
   - Enter the address (or start and end address) for the subnet you want to allow or deny.
   - If *Geo Location* is selected, select a geographic location or type of address.
6. Click *Save*.
7. Do one of the following:
   - Click *Global Settings* > *Access Control List* > *Access Control List.*
   - Click *Global Settings* > *Access Control List* > *Access Control List IPv6.*
8. Click *Add*.
9. Specify a name that identifies this item in the list.

10. For *Source address*, select the name you specified for the address or location in Step 3.
11. For *Action*, do one of the following:
    - If the type of the address you specified is *IP-Address*, *IPv6-Netmask*, or *IPv6-Range*, select *Deny* or *Track & Allow*.
    - If the type of the address you specified is *IP-Netmask* or *Geolocation* (IPv4 addresses only), select *Deny* or *Accept*.
12. Click *Save*.

To create IPv4 addresses or locations to add to the global ACL via the CLI

Enter the following commands:

```
config ddos global address
  edit <address_name>
    set type {ip-netmask | ip-address | geo-location}
    set ip-netmask <address_ipv4mask>
    set ip-address <address_ipv4>
    set geo-location <country_code>
end
```

where:

- `<address_name>` specifies the name that identifies the address in the UI
- `{ip-netmask | ip-address | geo-location}` specifies whether the address is an IP netmask, IP address, or geographic location
- `<address_ipv4mask>` specifies a subnet using either an Internet Protocol version 4 address and netmask separated by a space, or an address and CIDR-notation netmask separated by a slash
- `<address_ipv4>` specifies an Internet Protocol version 4 address
- `<country_code>` specifies a two-letter country code (ISO 3166-1-alpha-2 code);to specify Anonymous Proxy, enter `A1`; to specify Satellite Provider, enter `A2`

To create IPv6 addresses to add to the global ACL via the CLI

Enter the following commands:

```
config ddos global address-v6
  edit <address_name>
    set type {ipv6-network | ipv6-address}
    set ipv6-network <address_ipv6mask>
    set ipv6-address <address_ipv6>
end
```

where:

- `<address_name>` specifies the name that identifies the address in the list of addresses and ACL
- `{ipv6-network | ipv6-address}` specifies whether the address is an IP netmask or IP address
- `<address_ipv6mask>` specifies a subnet using an Internet Protocol version 6 address and netmask separated by a space
- `<address_ipv6>` specifies an Internet Protocol version 6 address

To add addresses or locations to the global ACL via the CLI

Enter the following commands:

```
config ddos global {acl | acl6} <acl_item_name> source-address
<address_name>
     set action {accept | deny}
end
```

where:

- `{acl | acl6}` specifies whether the address item is added to the IPv4 or IPv6 global ACL
- `<acl_item_name>` specifies the name that identifies the item in the IPv4 or IPv6 global ACL
- `<address_name>` specifies the name of the IPv4 or IPv6 address from the list of addresses that you created
- `{accept | deny}` specifies whether FortiDDoS allows packets that use the address (whitelist) or blocks them; you cannot specify `accept` for `ip-netmask` and `geo-location` addresses

See also

- Access and tracking control for Service Protection Profiles
- Creating an IP address item to use with ACLs
- Creating a service item to use with ACLs
- Add an address or service to a profile's ACL
- Do Not Track Policy list

Access and tracking control for Service Protection Profiles

The ACL for an SPP contains two types of items: IP addresses and services.

A service identifies packets using one of the following values:

- Fragment
- Protocol
- TCP Port
- UDP Port
- ICMP Type/Code
- URL
- HTTP header field: Host, Referer, Cookie, User Agent

Rules cannot combine multiple header values. For example, you can block packets from an IP address or block packets coming to a destination port. But you cannot block packets coming from that a specified IP address to a specified port.

Allow or deny protocols and ports

When you create a service that specifies protocols or TCP or UDP ports, you enter the start and end values for a range. To specify a single protocol or port, enter the same value for the start and end of the range.

Allow or deny URL

When you create a service that specifies a URL to allow or deny, enter the text that follows the protocol and the web address. For example, if you enter `http://www.website.com/index.html` in a browser to access a specific URL, enter `/index.html`.

Because the number of possible URLs is infinite, FortiDDoS stores these values in a hash table. Up to 32 767 such hash indexes are allowed. If there are duplicate hash-indexes, the most recent URL that corresponds to a hash index overwrites any previous URLs in the URL field. However, all the URLs affect the threshold and maximum packet rate calculations and all URLs that hash to the same index are denied if the hash index is blocked. Similarly, if there is an attack that corresponds to a hash index, all URLs that hash to the same location are dropped.

Allow or deny HTTP header field

You can accept or deny trafffic by specifying the following HTTP header field types: Host, Referer, Cookie, and User-Agent. This is useful when a specific hash-index is under attack. FortiDDoS allows the source to establish the TCP connection with the server. However, when FortiDDoS detects the specified hash-index, it denies the packet and sends an RST packet to the server to aggressively age the connection. The appliance treats all subsequent packets from the source on that TCP connection as foreign packets and blocks the source for the specified blocking period.

See also

- Add addresses or locations to the global ACL
- Creating an IP address item to use with ACLs
- Creating a service item to use with ACLs
- Add an address or service to a profile's ACL
- Blocking a protocol for a specified subnet

Creating an IP address item to use with ACLs

To create an IP address item to use with ACLs via the web UI

1. Do one of the following:
   - Click *Protection Profiles* > *Address* > *Address Config*.
   - Click *Protection Profiles* > *Address* > *Address Config IPv6*.
2. For *Service Protection Profiles*, select the profile you want to configure, and then click *Add*.
3. For *Name*, enter the name that you want to use to identify the address in the ACL.

   The name cannot contain spaces.
4. For *Address*, enter the name of the IP address to allow or deny.
5. Click *Save*.

To create an IP address item to use with ACLs via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp {address | address6} <address_name>
  spp-source-ip-address {<address_ipv4> | <address_ipv6}
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `{address | address6}` specifies whether the address is IPv4 or IPv6
- `<address_name>` specifies the specifies the name that identifies the address in the list of addresses
- `{<address_ipv4> | <address_ipv6}` specifies the IP address to allow or deny

See also

- Add addresses or locations to the global ACL
- Access and tracking control for Service Protection Profiles
- Creating a service item to use with ACLs
- Add an address or service to a profile's ACL

Creating a service item to use with ACLs

To create a service item to use with ACLs via the web UI

1. Click *Protection Profile* > *Service* > *Service Config*.
2. For *Service Protection Profiles*, select the profile you want to configure, and then click *Add*.
3. For *Name*, enter the name that you want to use to identify the service in the ACL.

   The name cannot contain spaces.
4. For *Type*, select one of the following options:
   - Fragment
   - Protocol
   - TCP-Port
   - UDP-Port
   - ICMP-Type-Code
   - URL
   - Host
   - Referer
   - Cookie
   - User-Agent
5. Specify additional information, if required.

   For protocols and TCP UDP ports, enter values that specify the start and end of a range.

6. Click *Save*.

Networks use of some of the protocols, such as 1 (ICMP), TCP (6), and UDP (17), ubiquitously. Ensure that you understand your network and its packet behavior before you use the ACLs for protocols.

Some Internet technologies, such as multimedia streaming, rely on fragmentation. Ensure that you understand your network and its packet behavior before you use the ACLs for fragmented packets.

To create a service item to use with ACLs via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp service
    edit <service_name>
    set type {fragment | protocol | tcp-port | udp-port |
    icmp-type-code | url | host | referer | cookie | user-agent}
    [set protocol-start <int_start>]
    [set protocol-end <int_end>]
    [set tcp-port-start <int_start>]
    [set tcp-port-end <int_end>]
    [set udp-port-start <int_start>]
    [set udp-port-end <int_end>]
    [set icmp-type <integer>[
    [set icmp-code <integer>]
    [set http-param <http_para_str>]
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `<service_name>` specifies the name that identifies the service in the user interface
- `{fragment | protocol | tcp-port | udp-port | icmp-type-code | url | host | referer | cookie | user-agent}` specifies the type of service
- `<int_start>` specifies the first value in a range of values, such as a range of protocols or TCP or UDP ports; for a single value, the start and end values are the same
- `<int_end>` specifies the first value in a range of values, such as a range of protocols or TCP or UDP ports
- `<integer>` specifies a value that defines the service, such as an ICMP type number or code
- `<http_para_string>` is a string that specifies an HTTP parameter, such as a URL or HTTP header field

See also

- Add addresses or locations to the global ACL
- Access and tracking control for Service Protection Profiles
- Creating an IP address item to use with ACLs
- Add an address or service to a profile's ACL

Add an address or service to a profile's ACL

To add an address or service to a profile's ACL via the web UI

1. Click *Protection Profiles* > *Access Control List* > *Access Control List*.
2. For *Service Protection Profiles*, select the profile you want to configure, and then click *Add*.
3. For *Name*, enter a name that identifies the rule in the ACL.
4. For *Type*, do one of the following:
   - Select *Address*, select the name of the address in the *Source Address* list, and then select *Deny* or *Track and allow*.
   - Select *Address IPv6*, select the name of the address in the *Source Address* list, and then select *Deny* or *Track and allow*.
   - Select *Service*, the traffic direction to control, and the name of the service in the *Service* list. Then, select *Deny* or *Accept*.
5. Click *Save*.

To add an address or service to a profile's ACL via the CLI

Enter the following commands:

```
edit <spp_name>
   config ddos spp acl type {v4address | service | v6address}
      set direction {outbound | inbound}
      set source-address <address_name>
      set service <service_name>
      set v6address <address_name>
      set service-action {accept | deny}
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `{v4address | service | v6address}` specifies the type of ACL item to add
- `{outbound | inbound}` specifies the traffic direction that the ACL item applies to
- `<address_name>` specifies the name of the IPv4 or IPv6 address from the list of addresses that you created
- `{accept | deny}` specifies whether FortiDDoS allows packets that use the address (whitelist) or blocks them

See also

- Add addresses or locations to the global ACL
- Access and tracking control for Service Protection Profiles
- Creating an IP address item to use with ACLs
- Creating a service item to use with ACLs

Blocking a protocol for a specified subnet

You cannot configure the ACL that blocks specific protocols for a specified subnet using the web UI.

To block a protocol for a specified subnet via the CLI

Enter the following commands:

```
config ddos global distress-acl
  edit <entry_index>
    set ip-netmask <address_ipv4mask>
    set protocol <protocol_int>
  next
  edit <entry_index>
    set ip-netmask <address_ipv4mask>
    set protocol <protocol_int>
end
```

where:

- `<entry_index>` is index number of the item in the ACL. You can enter a maximum number of 512 items.
- `<address_ipv4mask>` is a dotted decimal IPv4 address and netmask separated by a space, such as `192.168.1.99 255.255.255.0`, that specifies the subnet that the restriction applies to.
- `<protocol_int>` specifies the protocol that FortiDDoS blocks.

To restore access for all protocols that are blocked from specified subnets via the CLI

1. Enter the following commands:

```
config ddos global distress-acl
  purge
```

The following message is displayed.

```
This operation will clear all table!
Do you want to continue? (y/n)
```

2. Enter `y` (yes).

To restore access for a protocol to a specified subnet via the CLI

Enter the following commands:

```
config ddos global distress-acl
  edit <entry_index>
    unset ip-netmask
    unset protocol
end
```

where `<entry_index>` is index number of the item in the ACL.

Use the following commands to retrieve the index number of the ACL item:

```
config ddos global distress-acl
  show
```

See also

- Add addresses or locations to the global ACL
- Access and tracking control for Service Protection Profiles
- Creating an IP address item to use with ACLs
- Creating a service item to use with ACLs

FortiGuard IP Reputation Service

FortiGuard IP Reputation Service provides lists of IP addresses and network ID ranges that pose a threat to your network. After you register the serial number of your FortiDDoS appliance, you can schedule IP reputation list updates.

To download the lists, your appliance's management port requires access to the Internet. Alternatively, you can obtain the FortiGuard IP Reputation Service definitions file and upload it using the License Information widget.

You can specify how often the list is automatically updated and what types of reputation information to download. After you enable the feature, FortiDDoS downloads the most recent definitions file and then updates it according to the schedule you specified.

The dashboard's License Information widget displays the status of the most recent download attempt. If the download is successful and new definitions are available, the lists are replaced; otherwise, FortiDDoS maintains the previous versions.

To enable the IP Reputation Service via the web UI

1. Click *Global Settings* > *IP Reputation* > *IP Reputation*.
2. For Status, click *Enable*.
3. If you want your FortiDDoS appliance to connect to a Fortinet Distribution Server (FDS) other than the default one for its time zone, for *Override Server IP*, click *Enable*. Then, for *IP Address*, enter the appropriate IP address.

   You cannot choose a specific port on the override server. However, the tunneling option enables you to specify a port.
4. For *Schedule Type*, select how often the FortiGuard IP Reputation Service updates the reputation information.

   If you select *daily*, select the hour of the day when FortiDDoS updates reputation information. If you select *weekly*, select a day of the week for updates.
5. For *Category*, click the arrow beside the field, and then click categories to add or remove them from the field.

   Add the threat categories that are appropriate for your system. The *others* option adds protection against identified threats that do not fall into one of the current categories.
6. Optionally, for *Tunneling*, click *Enable*. Specify an IP address, port, and password for the tunneling server.

   When this option is enabled, FortiDDoS uses a tunnel to update definitions. Tunneling allows you to monitor the logs on the tunneling server to see when and how FortiDDoS sends update requests.
7. Click *Save*.

   The FortiDDoS appliance tests the connection to the Fortinet Distribution Network (FDN). If you specified an override server IP address, the appliance also tests that server. The time required to perform the tests depends on the speed of the appliance's network connection

and the number of timeouts that occur before the connection attempt is successful or the FortiDDoS appliance determines that it cannot connect.

FortiDDoS does not log events for these tests. However, if you configure tunneling, you can use the logs on the tunneling server to view information about connection attempts.

For more troubleshooting information, enter the following commands:

```
diagnose debug enable
diagnose debug application tp2ird
```

These commands display additional information in the CLI console. For example:

```
conf_process()
debug: conf_ddos_ip_reputation_set: 89
IPREn: 1, OverrideServerIP: 0, Schedule Freq: 2, IP Address:
0.0.0.0, proxyTunneling: 0, port: 443 proxyUser:
serial number from bios: FI800B3913800022
default server ip: 192.168.100.205
model name: FI800B
Firmware version: FI800B-FW-4.00-000
ServerOverrideIP=192.168.100.205, SerialNumber=FI800B3913800022,
FirmVersion=FI800B-FW-4.00-000, outdir=/var/intruguard/temp,
certdir=/tmp
create server list
create server list count : 4
Trying FDS 208.91.112.71:443 with AcceptDelta=1
Proxy tunneling is disabledcreating update packaging
Firmware Version: FI800B-FW-4.00-000
Serial Number : FI800B3913800022
Protocol=3.0|Command=Update|Firmware=FI800B-FW-4.00-000|SerialNumb
er=FI800B3913800022|UpdateMethod=0|AcceptDelta=1|DataItem=03000000
IRDB00201-00001.00015-1210151631*00000000FCNI00000-00000.00000-000
0000000*00000000FDNI00000-00000.00000-0000000000*01000000FSCI00100
-00000.00000-0000000000


Command:
Protocol=3.0|Command=Update|Firmware=FI800B-FW-4.00-000|SerialNumb
er=FI800B3913800022|UpdateMethod=0|AcceptDelta=1|DataItem=03000000
IRDB00201-00001.00015-1210151631*00000000FCNI00000-00000.00000-000
0000000*00000000FDNI00000-00000.00000-0000000000*01000000FSCI00100
-00000.00000-0000000000


Packing obj Type =0
Packing Identifier=FCPC
Packing Description=Command Object
Packing
obj=Protocol=3.0|Command=Update|Firmware=FI800B-FW-4.00-000|Serial
Number=FI800B3913800022|UpdateMethod=0|AcceptDelta=1|DataItem=0300
0000IRDB00201-00001.00015-1210151631*00000000FCNI00000-00000.00000
-0000000000*00000000FDNI00000-00000.00000-0000000000*01000000FSCI0
0100-00000.00000-0000000000
```

To enable the IP Reputation Service via the CLI

Enter the following commands:

```
config ddos global ip-reputation
  set ip-reputation-status {enable | disable}
  set override-server-ip {enable | disable}
  set ip-reputation-ip-address <override_server_address>
  set ip-reputation-schedule-type {hourly | daily | weekly}
  [set schedule-hour <hour_int>]
  [set schedule-weekdays {sunday | monday | tuesday ...|saturday}]
  set ip-reputation-category {phishing ddos anonymous-proxies spam
  others}
  set tunneling-status {enable | disable}
  set tunneling-address <tunneling_address>
  set tunneling-port <tunneling_port_int>
  set tunneling-username <tunneling_user_str>
  set tunneling-password <tunneling_pswd>
end
```

where:

- `{enable | disable}` specifies whether the setting is enabled
- `<override_server_address>` specifies the IP address of the Fortinet Distribution Server (FDS) to use for updates
- `<hour_int>` specifies at which hour in the day FortiDDoS updates the reputation lists when `ip-reputation-schedule-type` is hourly
- `{sunday | monday | tuesday ...|saturday}` specifies on the day of the week on which FortiDDoS updates the reputation lists when `ip-reputation-schedule-type` is weekly
- `{phishing ddos anonymous-proxies spam others}` specifies which types of reputation FortiDDoS dowloads
- `<tunneling_address>` specifies the IP address of the tunneling server
- `<tunneling_port_int>` specifies the port on the tunneling server to use
- `<tunneling_user_str>` specifies the username for the tunneling server
- `<tunneling_pswd>` specifies the password for the tunneling server

See also

- License Information widget

Enabling higher thresholds for proxy server IP addresses

FortiDDoS handles an IP address with the following characteristics as a legitimate proxy that is an intermediary for multiple IP addresses:

- Sources that have a consistently higher number of concurrent connections over a specified period of time when compared to a single-source IP address.
- Sources that provide the X-Forwarded-For or True-Client-IP HTTP header field.

To accommodate the normal behavior of a proxy IP address, FortiDDoS increases the thresholds that it applies to the proxy IPs addresses by a factor that you specify. It adjusts the

Most Active Source, SYN per source, and Concurrent Connections Per Source thresholds by multiplying the threshold for a standard IP address by the proxy IP threshold factor.

You set the factor by specifying the exponent of a power of two. For example, if you specify 5, the factor is 2^5 or 32. If the Most Active Source threshold is 1 000, and then the Most Active Source threshold that FortiDDoS applies to proxy IP addresses is 32 * 1000 or 32,000.

The maximum allowed threshold factor for proxy IP is 15 (2^{15} or 32,767).

Enable higher thresholds for proxy server IP addresses via the web UI

1. Click *Global Settings* > *Proxy IP* > *Proxy IP*.
2. To adjust thresholds for addresses that consistently have a high number of connections, click *Enable proxy IP auto detection,* and then complete the following settings:

Concurrent connections per source	Enter the minimum number of concurrent connections. When an IP address is the source of at least this number of connections during the specified observation period, FortiDDoS treats it as a proxy IP address.
Percent present	Enter the percentage of the proxy IPs that FortiDDoS adjusts. FortiDDoS adds this portion of the proxy IPs it detects to the proxy IP table and adjusts their thresholds. It adds addresses with the most connections first. For example, 100 proxy IPs have at least the number of connections specified by *Concurrent connections per source*. If *Percent present* is 30, FortiDDoS adjusts the thresholds for the 30 most active of these proxy IP addresses.
Past week	Specifies that FortiDDoS determines which IP addresses are proxies based on connections during the past week.
Past month	Specifies that FortiDDoS determines which IP addresses are proxies based on connections during the past month.

3. To adjust thresholds for addresses that provide the True-Client-IP or X-Forwarded-For HTTP header field, do the following:
   - Click *Enable proxy IP header detection*.
   - For *Proxy HTTP Header Type*, select *true-client-ip*, *X-Forwarded-For*, or both.
4. For *Proxy IP threshold factor*, specify the value that is used to adjust some thresholds for a proxy IP address as a factor of 2.
5. Click *Save*.

Enable higher thresholds for proxy server IP addresses via the CLI

Enter the following commands:

```
config ddos global proxy-ip-setting
  set auto-proxy-ip-status {enable | disable}
  set proxy-ip-percent-present <integer>
  set proxy-ip-observation-period {past-week | past-month}
  set header-proxy-ip-status {enable | disable}
  set header-proxy-type {true-client-ip X-Forwarded-For}
  set proxy-ip-threshold-factor <integer>
end
```

where:

- `{enable | disable}` specifies whether FortiDDoS uses the option to identify addresses to apply the proxy IP threshold factor to
- `<integer>` specifies a value for the setting
- `{past-week | past-month}` specifies the time period that FortiDDoS uses to identify IP addresses with a consistently high number of connections
- `{true-client-ip X-Forwarded-For}` specifies which HTTP header fields FortiDDoS uses to determine if an address is a proxy IP address

See also

- Viewing the current list of proxy IP addresses

Viewing the current list of proxy IP addresses

FortiDDoS allows you to view a list of the IP addresses that it has identified as a proxy IP.

This list is useful for identifying IP addresses that FortiDDoS is treating as a proxy but are actually attackers. You can add these kinds of IP addresses to an ACL to block their traffic. (See “Access control lists (ACLs)” on page 119.)

To view the current list of proxy IP addresses via the web UI

1. Click *Global Settings* > *Proxy IP* > *Proxy IP*.
2. Select *Backup proxy IP list*, and then click *Save.*

   Backup proxy IP list status displays the time and date when the list was last updated.
3. Click *Download* to view the list in your web browser.

To view the current list of proxy IP addresses via the CLI

Enter the following commands:

```
config ddos global proxy-ip-setting
  set download-list enable
  get backup-status <string>
end
```

where `<string>` is the time and date when the list was last updated

See also

- Enabling higher thresholds for proxy server IP addresses

Do Not Track Policy list

You can specify IP addresses that FortiDDoS does not restrict or track. You can add the following two types of IP addresses to the *Do Not Track Policy* list:

- **Do not track —** FortiDDoS never drops or blocks packets from these IP addresses, and does not include them in the statistics for continuous learning and threshold estimation.
- **Track and Allow—** FortiDDoS never drops or blocks packets from these IP addresses, and includes them in the statistics for continuous learning and threshold estimation.

You cannot add addresses specified by geolocation to the *Do Not Track Policy* list.

To add an address to the Do Not Track Policy list via the web UI

1. Create a named address (for instructions, see "Add addresses or locations to the global ACL" on page 122).
2. Do one of the following:
   - Click *Global Settings* > *Do Not Track Policy* > *Do Not Track Policy.*
   - Click *Global Settings* > *Do Not Track Policy* > *Do Not Track Policy IPv6.*
3. Click *Add*.
4. For *Name*, enter the name you want to use to identify this subnet in the policy list.
5. For *Source address*, select the name that you specified for the subnet address earlier.
6. For *Do not track action*, select *Do not track* or *Track and allow*, and then click *Save*.

To add an address to the Do Not Track Policy list via the CLI

Enter the following commands:

```
config ddos global {do-not-track-policy | do-not-track-policy-v6}
  edit <do_not_track_name>
    set do-not-track-source-address <no_track_ip_address>
    set do-not-track-action {track-and-allow | do-not-track}
  next
end
```

where:

- `{do-not-track-policy | do-not-track-policy-v6}` specifies whether the address is added to the IPv4 or IPv6 Do Not Track Policy list
- `<do_not_track_name>` specifies the name that identifies the item in the IPv4 or IPv6 Do Not Track Policy list
- `<no_track_ip_address>` specifies the IPv4 or IPv6 address
- `{track-and-allow | do-not-track}` specifies whether FortiDDoS tracks the source of traffic associated with the specified address

See also

- Add addresses or locations to the global ACL

SYN flood and zombie flood prevention

A SYN attack occurs when a target host is flooded with too many new TCP connection requests. Because TCP requires a three-way handshake to establish a connection, attackers that begin but do not finish the handshake process can absorb all resources reserved for legitimate users.

You can prevent SYN floods using several built-in techniques within FortiDDoS. The appliance provides the following resources for preventing SYN floods:

- SYN flood thresholds

  For each Service Protection Profile (SPP), you can create thresholds for incoming and outgoing traffic and on a per destination basis (corresponding to the most active destination).

  When the threshold for the maximum number of allowable SYN connection requests per second is exceeded, FortiDDoS initiates a SYN flood event. During this event, the appliance screens new connection requests to determine whether it should pass them to the protected system or block them.

- Legitimate IP (LIP) Address table

  FortiDDoS stores non-spoofed IP addresses that have done a three-way handshake successfully in a large table. Because this table retires entries every 5 minutes, it contains IP addresses that have recently connected successfully. Under SYN flood situation, that is, when the SYN flood threshold is crossed, FortiDDoS uses the LIP table to validate new connections. If the new connection request is from an address in this table, FortiDDoS allows it. It denies all other connections.

- *SYN flood mitigation mode*

  When this mode is enabled, FortiDDoS dynamically validates new connection requests to prevent spoofing by proxying the server to ensure that the client actually exists. If the client can respond to the FortiDDoS packets and they are valid, and then the client IP is added to the LIP table. For more information, see “Configuring SYN flood mitigation feature controls” on page 136.

- *new-connections* threshold

  To prevent zombie floods, you can set an appropriate threshold for legitimate IPs. This threshold ensures that legitimate IPs that have been compromised by attackers (also known as zombies) are not able to overwhelm the protected server. To set this threshold, add a Scalar threshold using *Protection Profiles* > *Thresholds* > *Thresholds*. For *Scalar Type*, select *new-connections*.

Other important scalar thresholds you adjust to prevent distributed DoS attacks:

- SYN packets per source (*syn*)
- Concurrent TCP connections per source (*concurrent-connections-per-source*)
- Concurrent TCP connections per destination (*concurrent-connections-per-destination*)

See also

- Configuring SYN flood mitigation feature controls
- Adjusting thresholds

Configuring SYN flood mitigation feature controls

For general information on SYN flood prevention, see “SYN flood and zombie flood prevention” on page 135.

To set SYN flood mitigation settings via the web UI

1. Click *Protection Profiles* > *SPP Settings* > *SPP Settings*.
2. For *Service Protection Profiles*, select the profile you want to configure.
3. For *SYN flood mitigation direction*, select *Inbound*, *Outbound*, or both.

4. Select one of the following modes:

Setting name	Description
SYN Cookie	Enables the anti-spoofing SYN proxy mechanism. SYN ACK is sent as part of the proxy mechanism. If the IP which sent the SYN responds with an ACK, a RST/ACK is generated and the IP is added to the legitimate IP address table. If the client then retries, it succeeds in making a TCP connection. This is the default value.
ACK Cookie	Enables an ACK-based anti-spoofing checking scheme. During a SYN flood attack, FortiDDoS sends the client two ACK packets: one with a correct ACK number and another with a wrong number. FortiDDoS determines whether the source is not spoofed based on the client's response. If the client's response indicates that the source is not spoofed, FortiDDoS allows the connection and adds the source to the legitimate IP address table.
SYN Retransmission	Enables a timeout-based SYN retransmission scheme. In this mode, FortiDDoS drops the initial SYNs and expects the client to send a SYN again to retry. If a preconfigured number of retransmitted SYNs arrive within a predefined time period, FortiDDoS considers the source to be legitimate. The appliance then allows the connection to go through and adds the source to the legitimate IP address table.

5. Click *Save*.

To set SYN flood mitigation settings via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp setting
    set syn-flood-mitigation-direction {inbound | outbound}
    set syn-flood-mitigation-mode {syn-cookie | ack-cookie |
    syn-retransmission}
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `{inbound | outbound}` specifies whether FortiDDoS verifies incoming or outgoing SYN packets for attack behavior
- `{syn-cookie | ack-cookie | syn-retransmission}` specifies how FortiDDoS determines if the source of SYN packets is legitimate.

See also

- SYN flood and zombie flood prevention

Configuring blocking periods

You can specify how long FortiDDoS blocks traffic that it has determined is part of an attack. By default, the FortiDDoS blocking period is 15 seconds (the maximum value). When blocking period is over, the threshold is checked again.

For example, if the value is 15, when FortiDDoS detects a flood, it blocks the packets associated with that flood for 15 seconds. If the traffic still qualifies as an attack after the blocking period expires, it blocks the traffic for another 15 seconds, and so on.

FortiDDoS also allows you to set a specific blocking period for IP addresses that it has identified as the source of an attack, which are called identified sources or source attackers. To identify source attackers, during the blocking period, FortiDDoS multiplies the packet rate from the source of the blocked packets by the value of *Source penalty factor inbound/outbound*. If the calculated rate exceeds the value of the *most-active-source* threshold, FortiDDoS identifies the IP address of the source as a source attacker.

You can specify a specific blocking period of up to 18 hours (approximately 64 000 seconds) for traffic from source attackers. The default blocking period for these sources is 60 seconds.

You can also configure an additional blocking period for sources that FortiDDoS identifies as the source of an attack and exceed a threshold for dropped packets that you specify. Like the blocking period for identified sources, you can specify a period of up to approximately 64 000 seconds and the default value is 60 seconds. By default, the threshold for the extended blocking period is 5,000 dropped packets.

To configure blocking periods via the web UI

1. Click *Global Settings* > *Settings* > *Settings*, and then configure these settings:

Setting name	Description
Blocking Period for all attacks (in seconds)	Enter a number between 1 and 15. Default is 0.
Blocking Period for Identified Sources (in seconds)	Enter a number between 1 and 65 535 seconds (approximately 18 hours). Default is 60.
Extended blocking Period for Identified Sources (in seconds)	Enter a number between 1 and 65 535 seconds (approximately 18 hours). Default is 60.
Drop Threshold to extend blocking period for Identified Sources	Specify the maximum number of dropped packets. When the number of dropped packets exceeds this threshold, FortiDDoS regulates the identified source using the extended blocking period.

2. Click *Save*.

To configure the blocking period via the CLI

Enter the following commands:

```
config ddos global setting
     set blocking-period <int>
     set source-blocking-period <int>
     set extended-blocking-period <int>
     set drop-threshold-within-blocking-period <int>
end
```

See also

- Setting penalty factors
- Effects of crossing a threshold
- Adjusting thresholds

Configuring the adaptive limit

The adaptive limit is an upper rate limit beyond which FortiDDoS blocks all traffic. It is the upper limit for the estimated threshold.

The adaptive limit is calculated as a percentage of the configured minimum threshold. For example, if the adaptive limit is 150 (the default), FortiDDoS can use its dynamic threshold estimation algorithm to raise the calculated threshold up to 150% of the value of the configured minimum threshold. Thus, if the inbound threshold for Protocols 17 (UDP) is 10,000, the threshold never falls below 10,000 and never exceeds 15,000.

When the adaptive limit is 100, when the configured threshold is reached, FortiDDoS makes no adjustment using dynamic threshold estimation.

In some cases, you want to keep the thresholds that you specified at all times. To ensure that the thresholds you set do not go above or below the specified values, set *Adaptive mode* to *Fixed*.

For more information about the adaptive limit, see "Continuous learning & adaptive threshold estimation" on page 23.

To adjust the adaptive limit via the web UI

1. Click *Protection Profiles* > *SPP Settings* > *SPP Settings*.
2. For *Adaptive limit (in percentage)*, enter a value from 100% to 300%.
3. Click *Save*.

To adjust the adaptive limit via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp setting
    set adaptive-mode adaptive
    set adaptive-limit <percent_int>
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `<percent_int>` is an integer that specifies a percentage of the configured minimum threshold

To disable adaptive thresholds via the web UI

1. Click *Protection Profiles* > *SPP Settings* > *SPP Settings*.
2. For *Adaptive mode*, select *Fixed*.
3. Click *Save*.

To disable adaptive thresholds via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp setting
    set adaptive-mode fixed
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)

See also

- Adaptive limit

Use the following steps to demonstrate and test how the FortiDDoS appliance blocks traffic. This test has the following requirements:

- A serial configuration like the one shown in Figure 25.
- The protected server responds to ICMP Echo (ping) packets.
- The connected, upstream system can generate a series of ICMP Echo Request packets.
- The Service Protection Profile used for the test is in prevention mode (see "Setting FortiDDoS to detection mode" on page 116). In detection mode, FortiDDoS reports the test attack but does not block packets.

Figure 25: Ping test configuration

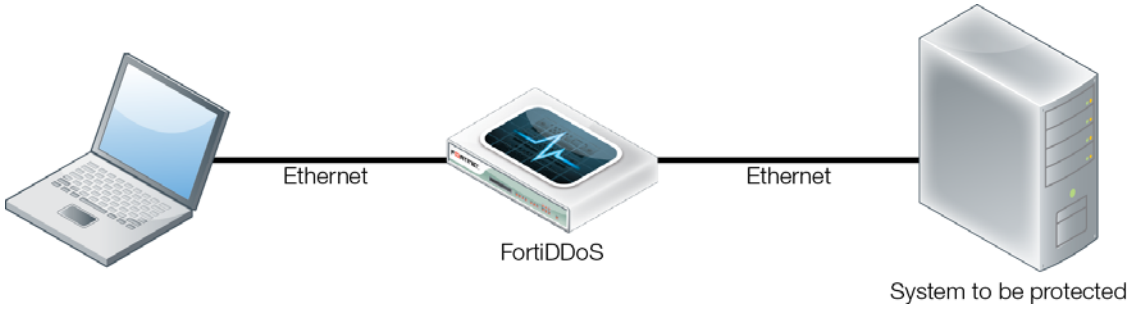

Configure the FortiDDoS appliance threshold for ping to 5 packets per second

1. Click *Protection Profiles* > *Thresholds* > *Thresholds*.
2. In the top-right corner of the content pane, for *Type*, select Protocols.
3. For *Service Protection Profiles*, select the profile you want to test.
4. Click *Add*.
5. In the New Thresholds dialog box, for *Name*, enter a name for the threshold.
6. For both *Protocol Start* and *Protocol End*, enter 1 (the protocol number for ICMP), and for both *Inbound Threshold* and *Outbound Threshold*, enter 5.
7. Click *Save*.

Generate ICMP (ping) traffic

From the PC/workstation, generate a small, controlled flood of 100 ICMP Echo (ping) packets directed to the protected system. For example, in UNIX/LINUX, the command line input looks like the following command. It generates an ICMP Echo Request (ping) packet to the specified address every 0.1 seconds until 100 packets are sent. This is the equivalent of 10 packets per second for 10 seconds:

```
ping -c 100 -i 0.1 AA.BB.CC.DD (where AA.BB.CC.DD represents the IP
address of the protected system)
```

The following screen capture is from an actual ping flood test. Notice that the first few pings are allowed to pass and receive a response.

As soon as the rate per second rises above the threshold, (somewhere in the first 11 packets) the FortiDDoS device blocks all ICMP packets for the 10-second threshold.

As soon as rate per second exceeds the threshold, FortiDDoS blocks the ICMP packets for one second, which returns the rate to 0. When the rate per second exceeds the threshold again, FortiDDoS blocks the packets for one second again, and so on.

In this sequence, this process is reflected by responses to the first 9 ping requests, followed by no response to packets 10 and 11 (which are blocked by the appliance). Then FortiDDoS allows packets 12-18 before the threshold is again reached.

```
[root@intel ISIC]# ping -c 100 -i 0.1 17.255.0.253
PING 17.255.0.253 (17.255.0.253) 56(84) bytes of data.
64 bytes from 17.255.0.253: icmp_req=1 ttl=64 time=0.109 ms
64 bytes from 17.255.0.253: icmp_req=2 ttl=64 time=0.137 ms
64 bytes from 17.255.0.253: icmp_req=3 ttl=64 time=0.133 ms
64 bytes from 17.255.0.253: icmp_req=4 ttl=64 time=0.129 ms
64 bytes from 17.255.0.253: icmp_req=5 ttl=64 time=0.132 ms
64 bytes from 17.255.0.253: icmp_req=6 ttl=64 time=0.131 ms
64 bytes from 17.255.0.253: icmp_req=7 ttl=64 time=0.131 ms
64 bytes from 17.255.0.253: icmp_req=8 ttl=64 time=0.132 ms
64 bytes from 17.255.0.253: icmp_req=9 ttl=64 time=0.132 ms
64 bytes from 17.255.0.253: icmp_req=12 ttl=64 time=0.132 ms
64 bytes from 17.255.0.253: icmp_req=13 ttl=64 time=0.115 ms
64 bytes from 17.255.0.253: icmp_req=14 ttl=64 time=0.132 ms
64 bytes from 17.255.0.253: icmp_req=15 ttl=64 time=0.130 ms
64 bytes from 17.255.0.253: icmp_req=16 ttl=64 time=0.131 ms
64 bytes from 17.255.0.253: icmp_req=17 ttl=64 time=0.138 ms
64 bytes from 17.255.0.253: icmp_req=18 ttl=64 time=0.129 ms
64 bytes from 17.255.0.253: icmp_req=22 ttl=64 time=0.136 ms
64 bytes from 17.255.0.253: icmp_req=23 ttl=64 time=0.133 ms
64 bytes from 17.255.0.253: icmp_req=24 ttl=64 time=0.131 ms
64 bytes from 17.255.0.253: icmp_req=25 ttl=64 time=0.129 ms
64 bytes from 17.255.0.253: icmp_req=26 ttl=64 time=0.132 ms
64 bytes from 17.255.0.253: icmp_req=27 ttl=64 time=0.111 ms
64 bytes from 17.255.0.253: icmp_req=28 ttl=64 time=0.129 ms
64 bytes from 17.255.0.253: icmp_req=32 ttl=64 time=0.134 ms
64 bytes from 17.255.0.253: icmp_req=33 ttl=64 time=0.130 ms
64 bytes from 17.255.0.253: icmp_req=34 ttl=64 time=0.133 ms
64 bytes from 17.255.0.253: icmp_req=35 ttl=64 time=0.129 ms
64 bytes from 17.255.0.253: icmp_req=36 ttl=64 time=0.132 ms
64 bytes from 17.255.0.253: icmp_req=37 ttl=64 time=0.130 ms
64 bytes from 17.255.0.253: icmp_req=38 ttl=64 time=0.133 ms
64 bytes from 17.255.0.253: icmp_req=42 ttl=64 time=0.142 ms
64 bytes from 17.255.0.253: icmp_req=43 ttl=64 time=0.134 ms
64 bytes from 17.255.0.253: icmp_req=44 ttl=64 time=0.141 ms
64 bytes from 17.255.0.253: icmp_req=45 ttl=64 time=0.137 ms
64 bytes from 17.255.0.253: icmp_req=46 ttl=64 time=0.139 ms
64 bytes from 17.255.0.253: icmp_req=47 ttl=64 time=0.133 ms
64 bytes from 17.255.0.253: icmp_req=48 ttl=64 time=0.132 ms
64 bytes from 17.255.0.253: icmp_req=52 ttl=64 time=0.135 ms
64 bytes from 17.255.0.253: icmp_req=53 ttl=64 time=0.128 ms
64 bytes from 17.255.0.253: icmp_req=54 ttl=64 time=0.131 ms
64 bytes from 17.255.0.253: icmp_req=55 ttl=64 time=0.129 ms
64 bytes from 17.255.0.253: icmp_req=56 ttl=64 time=0.134 ms
64 bytes from 17.255.0.253: icmp_req=57 ttl=64 time=0.132 ms
64 bytes from 17.255.0.253: icmp_req=58 ttl=64 time=0.130 ms
```

```
64 bytes from 17.255.0.253: icmp_req=62 ttl=64 time=0.135 ms
64 bytes from 17.255.0.253: icmp_req=63 ttl=64 time=0.131 ms
64 bytes from 17.255.0.253: icmp_req=64 ttl=64 time=0.133 ms
64 bytes from 17.255.0.253: icmp_req=65 ttl=64 time=0.129 ms
64 bytes from 17.255.0.253: icmp_req=66 ttl=64 time=0.130 ms
64 bytes from 17.255.0.253: icmp_req=67 ttl=64 time=0.131 ms
64 bytes from 17.255.0.253: icmp_req=68 ttl=64 time=0.131 ms
64 bytes from 17.255.0.253: icmp_req=71 ttl=64 time=0.134 ms
64 bytes from 17.255.0.253: icmp_req=72 ttl=64 time=0.133 ms
64 bytes from 17.255.0.253: icmp_req=73 ttl=64 time=0.129 ms
64 bytes from 17.255.0.253: icmp_req=74 ttl=64 time=0.131 ms
64 bytes from 17.255.0.253: icmp_req=75 ttl=64 time=0.114 ms
64 bytes from 17.255.0.253: icmp_req=76 ttl=64 time=0.114 ms
64 bytes from 17.255.0.253: icmp_req=77 ttl=64 time=0.114 ms
64 bytes from 17.255.0.253: icmp_req=81 ttl=64 time=0.114 ms
64 bytes from 17.255.0.253: icmp_req=82 ttl=64 time=0.115 ms
64 bytes from 17.255.0.253: icmp_req=83 ttl=64 time=0.114 ms
64 bytes from 17.255.0.253: icmp_req=84 ttl=64 time=0.113 ms
64 bytes from 17.255.0.253: icmp_req=85 ttl=64 time=0.114 ms
64 bytes from 17.255.0.253: icmp_req=86 ttl=64 time=0.114 ms
64 bytes from 17.255.0.253: icmp_req=87 ttl=64 time=0.115 ms
64 bytes from 17.255.0.253: icmp_req=91 ttl=64 time=0.118 ms
64 bytes from 17.255.0.253: icmp_req=92 ttl=64 time=0.114 ms
64 bytes from 17.255.0.253: icmp_req=93 ttl=64 time=0.113 ms
64 bytes from 17.255.0.253: icmp_req=94 ttl=64 time=0.115 ms
64 bytes from 17.255.0.253: icmp_req=95 ttl=64 time=0.114 ms
64 bytes from 17.255.0.253: icmp_req=96 ttl=64 time=0.114 ms
64 bytes from 17.255.0.253: icmp_req=97 ttl=64 time=0.113 ms

--- 17.255.0.253 ping statistics ---
100 packets transmitted, 72 received, 28% packet loss, time 10069ms
rtt min/avg/max/mdev = 0.109/0.127/0.142/0.011 ms
64 bytes from 172.16.0.50: icmp_seq=93 ttl=64 time=0.284 ms
```

The sequence states that 72 responses were received, which indicates that 28 packets were not received.

The number of blocked requests may vary depending on when the flood is started relative to the one-second blocking period.

See also

- Port Statistics graphs

Generating and reviewing a traffic statistics report

The *Traffic Statistics* option allows you to generate and view traffic statistics for a specific Service Protection Profile (SPP) and time period.

You can generate a traffic statistics report using the web UI or CLI. However, the web UI is required to view reports.

See also

- Generating a traffic statistics report
- Viewing a traffic statistics report
- Setting thresholds to system recommended values

Generating a traffic statistics report

FortiDDoS collects packets per second rates for all parameters during a 5-minute period. This rate is the rate of packets for a particular network parameter and Service Protection Profile (SPP) in both the inbound and outbound directions.

To generate a current traffic statistics report via the web UI

1. Click *Protection Profiles* > *Traffic Statistics* > *Generate*.
2. In the top-right corner of the content pane, for *Service Protection Profiles*, select the profile to configure thresholds for.
3. For *Period*, select a time period that corresponds to your learning period or other time period (for example, *1 week*).

   For *Status*, FortiDDoS displays either when it last generated a report for the profile and time period you specified or “Not available”.
4. If *Status* is not available, select *Generate*, click *Save*, and then confirm that you want to generate the report.

   FortiDDoS updates *Status* as it generates statistics for each traffic parameter. Click *Refresh Status* to see what is currently complete.

To generate a current traffic statistics report via the CLI

Enter the following commands:

```
edit <spp_name>
   config ddos spp threshold-report
      set generate {enable | disable}
      set report-period {last-hour | last-8-hours | last-24-hours |
      last-week | last-month | last-year}
end
```

where:

- `{enable | disable}` specifies whether to generate a report
- `{last-hour | last-8-hours | last-24-hours...}` specifies the time period for the traffic statistics report

See also

- Viewing a traffic statistics report

Viewing a traffic statistics report

After you generate a traffic statistics report, you can use *Traffic Statistics* > *Details* to view the maximum packet rate for each traffic parameter.

To view statistics that are reported using a single parameter and value (for example, SYN, SYN per source), from the *Type* list, select *Scalars*. Statistics that can be a list of parameters and values (for example, HTTP methods and protocols) are displayed individually in the *Type* list.

- Scalars (SYN, SYNPerSource, MostActiveSource, MostActiveDestination, SIPInvitePerSrc, SYNPerDst, ACKPerDst, RSTPerDst, FINPerDst)
- HTTP Methods (CONNECT, OPTIONS, HEAD, GET)
- Protocols
- TCP Ports
- UDP Ports
- ICMP Types/Codes
- URLs
- Hosts
- Referers
- Cookies
- User Agents

FortiDDoS does not display statistics that are lower than the following values:

Layer	Low threshold
3	100
4	500
7	200

To display values that are lower than these low threshold values, clear the *Do not show values below low threshold option*. (You cannot change the low threshold values that FortiDDoS applies to the traffic statistics details report.)

Traffic statistics details values are the maximum packets per second over the observation period (1 hour, 8 hours, 1 day, 1 week, 1 month, or 1 year). For example, during each 1-hour period, there are 12, 5-minute observation periods. FortiDDoS captures a maximum per second rate for each 5-minute interval. The maximum packets per second specifies the maximum value across these 12 periods of 5-minute intervals. However, the rate is still provided in packets per second.

For some items that are vectors, such as protocols, ports, and so on, the report lists the top 10 values. You can change the sort order as required.

FortiDDoS can use these statistics to generate thresholds for the *System Recommendation* option. For more information about *System Recommendation*, see “Setting thresholds to system recommended values” on page 146.

This maximum packet rate information is also available in the traffic statistics graphs. You do not have to manually generate the information to view it in a graph. See “Traffic graphs” on page 202.

To view a current traffic statistics report via the web UI

1. Click *Protection Profiles* > *Traffic Statistics* > *Details*.

2. Ensure that the Service Protection Profile (SPP) and time period that you generated statistics for are selected.
3. To display values that are lower than the low value threshold for the layer, clear *Do not show values below low threshold option*.
4. To review the results of the learning period, for *Type*, select each option and review the results.

   Alternatively, click *Save as PDF* to display information about all the available types in portable document format (PDF) file. Use the Adobe Reader tools to save or print the file for reference later.

See also

- Generating a traffic statistics report
- Setting thresholds to system recommended values

Setting thresholds to system recommended values

When it monitors and regulates subnets in prevention mode, FortiDDoS forwards packets until the traffic exceeds the threshold for a specific traffic parameter. When the traffic exceeds a threshold, it limits the offending traffic to the specified thresholds or blocks it for the configured blocking period. After blocking period, it checks the threshold again.

By default, FortiDDoS uses high thresholds so that the appliance passes all traffic without blocking.

You can adjust thresholds anytime and manipulate the granularity of each individual threshold in either direction for each Service Protection Profile. You can adjust individual thresholds or configure multiple thresholds at one time.

The *System Recommendation* option procedure sets most thresholds to values based on a traffic statistics report, which is the recommended method for setting thresholds for most types of traffic.

(UDP traffic requires its own service protection profile (SPP) and thresholds that are higher than the system recommended values. See “Creating an SPP for UDP traffic” on page 115.)

For information about additional ways of adjusting multiple thresholds, see “Adjusting multiple thresholds at one time” on page 152. For general threshold information, including information about setting individual thresholds, see “Adjusting thresholds” on page 150.

System Recommendation options

System Recommendation allows you to adjust all the thresholds associated with a specific layer as a group. You can adjust each group using the following methods:

- Set to factory defaults (high values)
- Set to a specified percentage of the value reported in a traffic statistics report
- Set a minimum threshold value for the layer

Preparing to use System Recommendation

To prepare to adjust thresholds using traffic history, use *Traffic Statistics* > *Generate* to update traffic statistics for the profile using an appropriate time period. Then, use *Traffic Statistics* >

Details to review the results. Note whether the maximums shown in the details are appropriate or should be adjusted.

Ensure that the traffic statistics report that you generate for use with *System Recommendation* is for a period that is free of attacks and that it is long enough to be a representative period of activity. You can use *Protection Profiles* > *Factory Reset* >*Factory Rest* set all thresholds to factory defaults to initiate a new learning period.

Factory Reset resets thresholds for a specific Service Protection Profile (SPP). It also resets other profile information such as attack data, attack graphs and traffic history.

Thresholds that are not set by System Recommendation

In most networks, some thresholds, such as the ones for the TCP and UDP protocols and some TCP ports, affect all traffic. For this reason, *System Recommendation* does not change the values for the following thresholds:

- TCP protocol (6) and UDP protocol (17)
- TCP Ports 21-23, 25, 53, 80, 110, 139, 443 and 590
- Most Active Destination, ACK Per Destination, RST Per Destination, FIN Per Destination, ESTAB Per Destination, Connection Per Destination

System Recommendation does not set high thresholds for UDP ports. Compared to TCP, there are only a limited number of ways to detect UDP attack traffic, such as UDP ports, the UDP protocol, and checking a source IP address against an ACL or legitimate IP address table. However, for TCP traffic, FortiDDoS can detect attacks at multiple layers, including new connections and concurrent connections on layer 4, and header fields, URLs, and HTTP methods on layer 7.

If required, you can change these thresholds individually using *Protection Profiles* > *Thresholds* > *Thresholds*. However, it is a best practice to detect attacks with the more specific thresholds for layer 7 traffic parameters or sources instead of service ports.

See also

- Set thresholds to system recommended values

Set thresholds to system recommended values

To set initial thresholds using System Recommendation via the web UI

1. After your initial learning period is complete (for example, after a week of operating a profile in detection mode), generate and review the profile's traffic statistics.

   Make note of any thresholds that you want to adjust during the threshold configuration process.

   For more information about traffic statistic reports, see "Generating and reviewing a traffic statistics report" on page 144.
2. Click *Protection Profiles* > *Thresholds* > *System Recommendation*.
3. In the top-right corner of the content pane, do the following:
   - For *Service Protection Profiles*, select the profile you want to configure.
   - For *Period*, specify the traffic statistics that FortiDDoS uses to calculate recommended thresholds by selecting a time period.

4. If there are no traffic statistics for the profile and time period that you selected, you are prompted to generate them using *Traffic Statistics* > *Generate*.

   If this message appears, generate the required traffic statistics and then return to this menu.

   When traffic statistics are available, the *System Recommendation* dialog box is displayed.

5. For each set of thresholds (grouped by OSI layer), select the following settings:

Setting name	Description
Percentage	Specifies that this group of thresholds is set using the specified percentage.
Factory defaults	Specifies that this group of thresholds is set to the factory default values instead of the recommended values. For more information about the factory default values, see "Set to factory defaults (high values)" on page 152.
Layer percentage	Specifies the values for this group of thresholds as a percentage of the incoming and outgoing maximums found in the traffic statistics. Enter a value between 100 and 300. For example, if the value is 100%, the threshold for traffic is the maximums in the report. If it is 300%, the rate threshold is three times the reported maximums.
Layer low traffic threshold	Specify the minimum threshold to use. If the recommended threshold calculated by FortiDDoS is lower than this value, the appliance uses this value as the threshold. For example, the maximum incoming packet rate for the layer 4 TCP protocol threshold in the traffic statistics report is 2,000 and the outgoing packet rate is 3,000. The value of *Layer 4 percentage* is 400 (percent) and the value of *Layer 4 low traffic threshold* is 10,000. FortiDDoS uses these settings to calculate a recommended inbound TCP protocol threshold of 8,000 (2,000 * 400% = 8,000). However, because 8,000 is less than the low traffic threshold of 10,000, the appliance uses 10,000 as the incoming threshold. The recommended threshold for outbound packets is 12,000 (3,000 * 400% = 12,000). Because 12,000 is greater than the low traffic threshold of 10,000, the appliance uses 12,000 as the outgoing threshold. This setting is helpful when you think that the traffic statistics maximums are too low to generate a useful threshold.

6. Click *Save*.

To set initial thresholds using System Recommendation via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp threshold-adjust
    set threshold-adjustment-type system-recommendation
    set threshold-system-recommended-report-period
      {1-hour | 8-hours | 1-day | 1-week | 1-month | 1-year}
    set threshold-system-recommended-layer-3 {layer3-percentage |
    layer3-factory-defaults}
    set threshold-system-recommended-layer-3-percentage
    <percent_int>
    set threshold-system-recommended-layer-3-low-traffic <integer>
    set threshold-system-recommended-layer-4 {layer4-percentage |
    layer4-factory-defaults}
    set threshold-system-recommended-layer-4-percentage
    <percent_int>
    set threshold-system-recommended-layer-4-low-traffic <integer>
    set threshold-system-recommended-layer-7 {layer7-percentage |
    layer7-factory-defaults}
    set threshold-system-recommended-layer-7-percentage
    <percent_int>
    set threshold-system-recommended-layer-7-low-traffic <integer>
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `{1-hour | 8-hours | 1-day...}` specifies the time period of the traffic statistics report that you generated
- `{layer3-percentage | layer3-factory-defaults}`,`{layer4-percentage | layer4-factory-defaults}`, and `{layer7-percentage | layer7-factory-defaults}` specify whether the thresholds for the layer are set to factory defaults (high values) or a percentage of the reported values
- `<percent_int>` specifies the value for thresholds for the layer as a percentage of the reported values
- `<integer>` specifies a minumum threshold value for thresholds for the layer

See also

- Setting thresholds to system recommended values

Monitoring attack statistics

You can view FortiDDoS monitoring activities using the web UI dashboards, traffic graphs, and reports.

- Dashboards

  Use the following dashboards to view events related to FortiDDoS monitoring and prevention activity:

  - *System* > *Status* > *Dashboard*
  - *Log & Report* > *Log Access* > *DDoS Attack Log*
  - *Log & Report* > *Report Browse* > *Executive Summary*
  - *Log & Report* > *Attack Graphs* > *Attack Graphs*
- Traffic graphs

  Graphs display information about traffic for individual Service Protection Profiles.

  For example, the ping test in "Testing your installation" on page 141 generates graph information, including information in the following graphs:

  - *Monitor* > *Port Statistics* > *Packets*
  - *Monitor* > *Specific Graphs* > Protocols
  - *Monitor* > *Specific Graphs* > *ICMP Types/Codes.*
- Reports

  You can create reports that summarize past attacks and use the Executive Summary and Attack Graphs dashboards to view current report information.

  For example, the ping test in "Testing your installation" on page 141 generates activity that is displayed in the following reports:

  - Top Attacked Services and Top Attacked ICMP Type and Code
  - Top Attacked Protocols
  - Top Attacks

For more information on these tools, see "Monitoring attack activity and other system information" on page 195.

Adjusting thresholds

FortiDDoS is a network behavior analysis (NBA) system. It allows you to set thresholds for all continuously monitored traffic. Each Service Protection Profile has its own set of thresholds, and you can set separate values for incoming and outgoing traffic for each threshold.

You can adjust thresholds anytime. FortiDDoS allows you to change the value of individual thresholds for each Service Protection Profile, for incoming taffic, outgoing traffic, or both.

By default, FortiDDoS uses high thresholds that pass all traffic without blocking.

In most cases, you start adjusting the thresholds after an initial learning period. (See "Setting FortiDDoS to detection mode" on page 116.)

This adjusting process can be interactive, based on the system's behavior. If you think that there are false positives (that is, the appliance has dropped or blocked legitimate traffic), you can increase the threshold values. If attacks go undetected, you can reduce the thresholds.

In most cases, after you have adjusted the thresholds, no further adjustments are required because FortiDDoS is continuously learning the traffic characteristics. For example, if *Adaptive*

mode is enabled, FortiDDoS increases the thresholds if it detects a general upward trend or seasonality (a predictable or expected variation) in the traffic levels.

However, threshold adjustment is required whenever the traffic that a profile protects changes thoroughly (for example, when you add a new service or replace a server).

Choosing threshold values

Setting the right thresholds is key to performance of a network behavior analysis (NBA) system. Because there are so many methods for controlling thresholds in FortiDDoS, it is not easy to guess correct thresholds.

FortiDDoS has features that can help you to set the right thresholds. For example, you can use the estimated thresholds based on the past history of traffic as the baseline (*Protection Profiles* > *Thresholds* > *System Recommendation*), if the time period was free of attacks. Alternatively, you can use the traffic graphs to determine a rate that you think is appropriate.

Avoiding disruptions while adjusting thresholds

To avoid disruptions while adjusting the thresholds, keep the Service Protection Profile that you are adjusting in detection mode. In this mode, FortiDDoS reports all attacks and identifies packets to drop but does not drop any packets.

Use detection mode when you adjust the thresholds until you stop seeing false positives. The dashboard's *SPP Attacks* widget provides a current count of dropped packets.

Once you are satisfied with the results of the mitigation activity, you can select prevention mode.

For more information on detection and prevention modes, see "Setting FortiDDoS to detection mode" on page 116.

For spoofed attacks, has a mechanism that proxies the requesting server to ensure that the client actually exists.**See also**

- Adjusting multiple thresholds at one time
- Adjusting thresholds individually

Adjusting multiple thresholds at one time

In addition to adjusting individual thresholds, FortiDDoS allows you to use the following methods to set multiple thresholds for each Service Protection Profile:

- Factory Defaults

  Sets the thresholds for a Service Protection Profile to factory defaults, which is the line rate value. The appliance passes all traffic without blocking any packets.

- Percent Adjust

  Adjust the minimum thresholds up or down by a percentage that you specify. This is useful when you expect a sudden flood of traffic (for example, because of a news story or an advertising campaign).

- Emergency Setup

  FortiDDoS adjusts only certain key thresholds based on empirical knowledge. You can expect these adjustments to protect against common attacks.

  This option is useful when you deploy the appliance in an unknown environment and do not have time for an initial learning period.

  It is recommended that you use this method only during the immediate attack. When the attack is over, use the normal process to set initial thresholds. See “Setting thresholds to system recommended values” on page 146.

- System Recommendation

  The most commonly used and recommended way to set the thresholds. The system recommended values are based on traffic statistics generated as part of the learning period. See “Setting thresholds to system recommended values” on page 146.

See also

- Setting thresholds to system recommended values
- Set to factory defaults (high values)
- Set to a percentage of current thresholds
- Set using Emergency Setup
- Adjusting thresholds individually

Set to factory defaults (high values)

In some situations, you may want to set all thresholds to factory defaults. For example:

- You have installed a new server or network behind FortiDDoS. Therefore you want the appliance to learn the traffic pattern without dropping the packets.
- You want to ensure that the application does not drop any packets due to rate thresholds.

Setting all thresholds to factory defaults does not set the appliance to factory defaults. To reset all configuration for a profile, see “Resetting profile data or the appliance configuration” on page 266.

To set the thresholds for just a single OSI layer (3, 4, or 7) to factory defaults, see “Setting thresholds to system recommended values” on page 146

To set all thresholds to factory defaults or high values via the web UI

1. Click *Protection Profiles* > *Thresholds* > *Factory Defaults*.

2. In the top-right corner of the content pane, for *Service Protection Profiles*, select the profile you want to configure.
3. Select *Set to factory defaults*.
4. Click *Save*.

   All thresholds for the profile are set to high values.

To set all thresholds to factory defaults or high values via the CLI

Enter the following commands:

```
edit <spp_name>
   config ddos spp threshold-adjust
      set threshold-adjustment-type factory-defaults
      set threshold-factory-defaults {enable | disable}
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `{enable | disable}` specifies whether FortiDDoS sets the profile thresholds to factory defaults

Set to a percentage of current thresholds

You can use the Percent Adjust options to adjust FortiDDoS for a one-time anticipated change in traffic.

This adjustment ensures that legitimate traffic is not blocked because of an increase in traffic that you expected. For example, when a news flash or other important announcement increases traffic to a company's web site.

For example, when you enter 10, all thresholds in the system are recalculated using the following formula:

New Minimum Threshold = Current Minimum Threshold + (Current Minimum Threshold * 10)/100

To return the thresholds to their original values, divide the percentage you used to adjust the threshold by 100. For example, if you changed the threshold to 120% of its original value, entering 83.33% returns it to a value close to the original.

To set thresholds to a percentage of current thresholds using the web UI

1. Click *Protection Profiles* > *Thresholds* > *Percent Adjust*.
2. In the top-right corner of the content pane, for *Service Protection Profiles*, select the profile you want to configure.
3. For *Adjust minimum threshold by percentage*, enter a value between 0 and 300.

   For example, to reduce the threshold by 20 per cent, enter 80. To increase it by 20 per cent, enter 120.
4. Click *Save*.

To set thresholds to a percentage of current thresholds via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp threshold-adjust
    set threshold-adjustment-type percent-adjust
    set threshold-percent-adjust <percent_int>
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `<percent_int>` specifies a new value for all thresholds as a percentage of the current thresholds

Set using Emergency Setup

This option adjusts only certain key thresholds based on empirical knowledge. You can expect these adjustments to protect against common attacks.

For example, use *Emergency Setup* if you are already under attack and need to deploy the unit without an initial learning period.

To set the Emergency Setup thresholds via the web UI

1. In the main menu, go to *Protection Profiles* > *Thresholds* > *Emergency Setup*.
2. In the top-right corner of the content pane, for *Service Protection Profiles*, select the profile you want to configure.
3. Edit the default values, if required
4. Click *Save*.

To set the Emergency Setup thresholds via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp threshold-adjust
    set threshold-adjustment-type easy-setup
    set threshold-easy-setup-inbound-syn-threshold <integer>
    set threshold-easy-setup-outbound-syn-threshold <integer>
    set threshold-easy-setup-inbound-syn-per-source-threshold
    <integer>
    set threshold-easy-setup-outbound-syn-per-source-threshold
    <integer>
    set threshold-easy-setup-inbound-most-active-source-threshold
    <integer>
    set threshold-easy-setup-outbound-most-active-source-threshold
    <integer>
    set
    threshold-easy-setup-inbound-concurrent-connections-per-source-
    threshold <integer>
    set
    threshold-easy-setup-outbound-concurrent-connections-per-source
    -threshold <integer>
    set
    threshold-easy-setup-inbound-concurrent-connections-per-destina
    tion-threshold <integer>
    set
    threshold-easy-setup-outbound-concurrent-connections-per-destin
    ation-threshold <integer>
    set
    threshold-easy-setup-inbound-concurrent-invite-per-source-thres
    hold <integer>
    set
    threshold-easy-setup-outbound-concurrent-invite-per-source-thre
    shold <integer>
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `<integer>` specifies a threshold

Adjusting thresholds individually

To specify thresholds that you set for a single parameter (for example, SYN, SYN per source), from the *Type* list, select *Scalars*. Thresholds that set maximum values for parameters that are a collection or range (for example, HTTP methods and protcols) are displayed individually in the *Type* list.

You can set the following thresholds:

Table 9: FortiDDoS configurable traffic thresholds

Name	Layer	Description
Scalars		
syn	4	There is a certain number of SYN packets that are normal in any given network. If that number exceeds a threshold, it could be due a SYN flood attack. If the number of the SYN packets in a Service Protection Profile (SPP) and in a particular direction exceeds the set SYN threshold, FortiDDoS goes into a SYN prevention mode. In this mode, FortiDDoS validates whether the IP is spoofed or not. It validates the addresses using the Legitimate IP address table, which it fills during normal operations with IP address entries that complete a 3-way TCP handshake. A SYN flood also triggers the SYN flood mitigation mode, which validates IP addresses that are not already in the Legitimate IP address Table.
syn-per-src	4	There is a certain rate of SYN packets that is normal for any source to send. If the rate exceeds a threshold, it could be a deliberate attack. This threshold blocks any single source from launching a TCP SYN attack and blocks it from establishing new connections for a certain duration.
most-active-source	3	Every packet comes from a specific source. FortiDDoS calculates the traffic rates for the most active source on a per second basis and allows you to set threshold. No single source can exceed this limit. FortiDDoS can track up to 1 million IP source addresses simultaneously.
concurrent-connections-per-source	4	There is a certain number of TCP connections that a single source builds. If that number exceeds a threshold, it could be a deliberate attack. This threshold blocks any single source from launching slow or fast TCP connection attacks and blocks it from establishing new connections for a certain duration. In addition, FortiDDoS sends TCP RST packets to the server to aggressively age the connections from the server. This process ages all existing connections from the FortiDDoS and the server.

Table 9: FortiDDoS configurable traffic thresholds (continued)

Name	Layer	Description
most-active-destination	3	Every packet is addressed to a specific destination. FortiDDoS calculates the traffic rates for most active destination on a per second basis and allows you to set a threshold. No single destination can exceed this limit. FortiDDoS can track up to 1 million IP destinations simultaneously.
fragment	3	Although the IP specification allows IP fragmentation, an excessive number of fragmented packets can cause some systems to hang or crash. The threshold specifies the number of acceptable fragmented packets per second.
sip-invite-per-src	7	This thresholds applies to SIP (Session Initiation Protocol) packets sent over either TCP or UDP and blocks any single source (for configured blocking period) from launching a SIP INVITE flood attack. There is a certain rate of SIP INVITE requests that is normal for any source to send. A rate that exceeds this threshold could be a deliberate attack.
sip-concurrent-invite-per-src	7	This thresholds applies to SIP (Session Initiation Protocol) packets sent over either TCP or UDP and blocks any single source (for configured blocking period) from launching a SIP INVITE flood attack. There is a certain rate of SIP INVITE requests that is normal for any source to send at one time. A rate that exceeds this threshold could be a deliberate attack.
sip-register-per-src	7	This thresholds applies to SIP (Session Initiation Protocol) packets sent over either TCP or UDP and blocks any single source (for configured blocking period) from launching a SIP REGISTER flood attack. There is a certain rate of SIP REGISTER requests that is normal for any source to send. A rate that exceeds this threshold could be a deliberate attack.

Table 9: FortiDDoS configurable traffic thresholds (continued)

Name	Layer	Description
new-connections	4	FortiDDoS tags IP addresses that perform a three-way-handshake as legitimate IP addresses. It continuously learns the normal rate for three-way-handshakes. When a SYN flood is detected, the FortiDDoS examines the IP source and IP destination addresses of all incoming SYN packets to determine if the addresses are known to be good. The number of connection requests from these “legitimate IP” sources is counted each second. If the rate of new connections in any second exceeds the *new-connections* threshold, the system reports a zombie flood to indicate that sources that were formerly legitimate are now (probably unwitting) attackers. When this happens, FortiDDoS blocks all new connection requests for the configurable blocking period. The *new-connections* threshold should always be higher than the SYN Packet threshold in order to be effective. It is recommended that you use the FortiDDoS generated threshold unless you have a specific reason to change it.
syn-per-dst	4	Specifies the maximum rate for SYN packets to a single destination.
ack-per-dst	4	Specifies the maximum rate for ACK packets to a single destination.
rst-per-dst	4	Specifies the maximum rate for RST packets to a single destination.
fin-per-dst	4	Specifies the maximum rate for FIN packets to a single destination.
estab-per-dst	4	Specifies the maximum number of established TCP connections for a single destination.
http-urls-per-source	7	Specifies the maximum number of HTTP accesses that a source can perform during a specified observation period. To specify the observation period, go to *Global Settings* > *Settings* > *Settings*.

Table 9: FortiDDoS configurable traffic thresholds (continued)

Name	Layer	Description
http-sequential-access-per-source	7	Specifies the maximum number of times a single IP address can retrieve the same URL back-to-back (that is, without accessing other URLs in between) within a specified observation period. This behavior is often evidence of a scripted attack, where bots and not humans perform the activity.
mandatory-header-violation	7	The *mandatory-header-violation* threshold specifies the maximum number of packets that do not have the mandatory HTTP header fields that a single IP can send. If the source exceeds this maximum, FortiDDoS identifies it as an offender. HTTP header fields contain the operating parameters of an HTTP request or response. With the request or response line (the first line of the message), they form the message header. The header fields define various characteristics of the requested data transfer or the data that is provided in the message body. If a packet does not have all the mandatory headers, it is often evidence of a scripted attack. You can also specify the minimum number of HTTP header fields that are required in a GET access to a URL. For more information, see "Set the mandatory HTTP header count" on page 179.
concurrent-connections-per-destination	4	There is a certain number of TCP connections that a single destination can have during normal times. If that number exceeds a threshold, it could be a deliberate attack. This threshold blocks any single destination from getting slow or fast TCP connection attacks. FortiDDoS blocks the destination from establishing new connections for a certain duration. In addition, FortiDDoS sends TCP RST packets to the destination to aggressively age idle connections from the destination server. This process ages all idle connections from the FortiDDoS and the server.

Table 9: FortiDDoS configurable traffic thresholds (continued)

Name	Layer	Description
Other thresholds		
HTTP Methods	7	HTTP/1.1 uses the following set of common methods: • GET • HEAD • OPTIONS • TRACE • POST • PUT • DELETE • CONNECT Each method has a normal behavior in a given website in terms of number of packets per second. When an abnormal attack happens, these normal thresholds are breached.
Protocols	3	The IP header contains a Protocol field that specifies the upper layer protocol. FortiDDoS continuously learns traffic rates for each protocol and allows you to set thresholds for each of them independently.
TCP Ports	4	FortiDDoS allows you to specify a packet-per-second limit for a TCP Port number. This is helpful to prevent floods against a specific application such as HTML, FTP, SMTP or SQL. TCP accommodates 64K (65,536) ports, most of which may never be used by a particular server. Conversely, a server may see most or all of its traffic on a small group of TCP ports. For this reason, globally assigning a single threshold to all ports generally does not provide useful protection. However, you can globally set a (usually low) TCP Port Threshold for all TCP ports and then manually configure a higher threshold for the ports your protected system is using.
UDP Ports	4	The User Datagram Protocol (UDP) Port threshold limits the number of packets that are sent to or by a specific UDP port each second. If the threshold for a particular port number is reached, all connections that port sends or is sent are blocked for the configured blocking period.

Table 9: FortiDDoS configurable traffic thresholds (continued)

Name	Layer	Description
ICMP Types/Codes	4	The header of Internet Control Message Protocol packets include an 8-bit type field, followed by an 8-bit code field. The value of this field can be read as a hexadecimal number. The ICMP Threshold determines the number of packets-per-second that are allowed by packets using specific ICMP type and code values. If the threshold is reached for a particular type and code combination, all ICMP packets containing that type/code combination are blocked for the configured blocking period.
URLs	7	The GET or POST operations in HTTP involve URLs. For example, if you access http://www.website.com/index.html, the URL is /index.html. Botnets make it easy to launch attacks on specific URLs. When such an attack happens, FortiDDoS can isolate the URL and limit just the traffic that is associated with it, while all other traffic is unaffected. FortiDDoS allows you to set the packets per second threshold for up to 32,767 hash indexes per Service Protection Profile (SPP) in two directions. FortiDDoS dynamically hashes the URLs that you specify into a bucket of up to 32,767 indexes and adjusts the corresponding thresholds adaptively. This ensures that thresholds are set for a specific website and profile.

Table 9: FortiDDoS configurable traffic thresholds (continued)

Name	Layer	Description
Hosts	7	HTTP GET operations involve some HTTP header fields. Four of them are important for flood identification. With the advent of botnets, it is easy to launch attacks using scripts. Most of the scripts use the same code. The chances that they all use the same Host, Referer, Cookie, or User-Agent header fields is very high. When such an attack happens, FortiDDoS can easily isolate the four headers among many and limit traffic associated with that specific header, while all other traffic is unaffected. FortiDDoS allows you to set the threshold for up to 512 header hash indexes per Service Protection Profile (SPP) in two directions.
Referers	7	
Cookies	7	
User Agents	7	
SIP User Agents	7	Like HTTP, SIP can have a User-Agent header field. In many cases, a scripted attack generates packets with identical User-Agent header fields. FortiDDoS can easily isolate this header and limit traffic associated with it without affecting any other traffic. FortiDDoS allows you to set the threshold for up to 512 header hash indexes per Service Protection Profile (SPP) in two directions.

To set an individual threshold via the web UI

1. Click *Protection Profiles* > *Thresholds* > *Thresholds*.
2. For *Type*, select the type of threshold you want to set.
3. For *Service Protection Profile*, select the profile you want to add the threshold to.
4. Click *Add*, and then specify the name, the inbound and outbound threshold and any other values required for the threshold.

   For example, for *Scalar* thresholds, select a *Scalar Type* value. For *ICMP Types/Codes*, specify an ICMP type and code.
5. Click *Save*.

To set an individual scalar threshold via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp scalar-threshold
    edit <threshold_name>
    set type {syn | syn-per-src | most-active-source |
    concurrent-connections-per-source | most-active-destination |
    fragment | sip-invite-per-src | sip-concurrent-invite-per-src |
    sip-register-per-src | new-connections | syn-per-dst |
    ack-per-dst | rst-per-dst | fin-per-dst | estab-per-dst |
```

```
    http-urls-per-source | http-sequential-access-per-source |
    mandatory-header-violation}
    set inbound-threshold <integer>
    set outbound-threshold <integer>
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `{syn | syn-per-src | most-active-source...}` specifies the type of threshold
- `<threshold_name>` specifies the name that identifies the threshold in the list of thresholds
- `<integer>` specifies the value of the incoming or outgoing threshold

To set an individual HTTP method threshold via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp http-method-threshold
    edit <threshold_name>
    set method {get | head | options | trace | post | put | delete |
    connect}
    set inbound-threshold <integer>
    set outbound-threshold <integer>
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `<threshold_name>` specifies the name that identifies the threshold in the UI
- `{get | head | options | trace | post | put | delete | connect}` specifies the type of threshold
- `<integer>` specifies the value of the incoming or outgoing threshold

To set an individual SIP method threshold via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp sip-method-threshold
    edit <threshold_name>
    set method {invite | register | refer | options}
    set inbound-threshold <integer>
    set outbound-threshold <integer>
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `<threshold_name>` specifies the name that identifies the threshold in the list of thresholds
- `{invite | register | refer | options}` specifies the type of threshold
- `<integer>` specifies the value of the incoming or outgoing threshold

To set an individual protocol threshold via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp protocol-threshold
    edit <threshold_name>
    set protocol-start <protocol_int>
    set protocol-end <protocol_int>
    set inbound-threshold <integer>
    set outbound-threshold <integer>
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `<threshold_name>` specifies the name that identifies the threshold in the list of thresholds
- `<protocol_int>` specifies the start or end value for a protocol range (for a single port, specify the same value for `set protocol-start` and `set protocol-end`)
- `<integer>` specifies the value of the incoming or outgoing threshold

To set an individual TCP or UDP port threshold via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp {tcp-port-threshold | udp-port-threshold}
    edit <threshold_name>
    set port-start <port_int>
    set port-end <port_int>
    set inbound-threshold <integer>
    set outbound-threshold <integer>
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `<threshold_name>` specifies the name that identifies the threshold in the list of thresholds
- `<port_int>` specifies the start or end value for a port range (for a single port, specify the same value for `set port-start` and `set port-end`)
- `<integer>` specifies the value of the incoming or outgoing threshold

To set an individual ICMP type and code threshold via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp icmp-type-code-threshold
    edit <threshold_name>
    set icmp-type-code-start <type_code_int>
    set icmp-type-code-end <type_code_int>
    set inbound-threshold <integer>
    set outbound-threshold <integer>
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `<threshold_name>` specifies the name that identifies the threshold in the list of thresholds
- `<type_code_int>` specifies the start or end value for a ICMP type and code range; calulate the value using the following formula: 256 * ICMP type number + ICMP code; for a single ICMP type and code, specify the same value for `set icmp-type-code-start` and `set icmp-type-code-end`
- `<integer>` specifies the value of the incoming or outgoing threshold

To set an individual URL or HTTP header field threshold via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp {http-url-threshold | http-host-threshold |
  http-referer-threshold | http-cookie-threshold |
  http-user-agent-threshold}
    edit <threshold_name>
    set {url | host | referer | cookie | user-agent} <string>
    set inbound-threshold <integer>
    set outbound-threshold <integer>
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `{http-url-threshold | http-host-threshold | http-referer-threshold ...}` specifies the type of threshold to set
- `<threshold_name>` specifies the name that identifies the threshold in the list of thresholds
- `<string>` specifies the value for the URL or HTTP header field
- `<integer>` specifies the value of the incoming or outgoing threshold

See also

- Adjusting multiple thresholds at one time
- Specifying Protocols, TCP Ports, and UDP Ports thresholds
- Index numbers for URLs and HTTP header fields
- ICMP type/code threshold and “Echo groping”

Specifying Protocols, TCP Ports, and UDP Ports thresholds

When you specify a threshold for protocols, TCP ports, and UDP ports, you enter a range, even if you are specifying a threshold for a single protocol or port.

For example, to set a threshold for protocol 6, for both *Protocol Start* and *Protocol End*, enter 6.

See also

- Adjusting multiple thresholds at one time

Index numbers for URLs and HTTP header fields

When you specify a threshold for a URL or HTTP header field (Host, Referer, Cookie, or User-Agent), FortiDDoS generates a corresponding hash value. You can use this value to specify the threshold elsewhere in the web UI.

For example, when you create a threshold for a URL, FortiDDoS displays its hash index in the list of URL thresholds. To view statistics associated with the threshold, click *Monitor* > *Specific Graphs* > *URLs*, and then, for *Please enter URL/Hash index*, enter either the original URL you specified or the hash.

See also

- Adjusting multiple thresholds at one time
- Specific Graphs

ICMP type/code threshold and “Echo groping”

A popular use for ICMP is the “Echo groping” message (type 8) and its corresponding reply (type 0), which are often useful tools to test connectivity and response time. In some cases, this message and reply can also be used as an attack weapon to effectively disable a target system’s network software. Take care when you set the ICMP type 0 and type 8 thresholds to ensure the desired functionality is preserved.

See also

- Adjusting multiple thresholds at one time

Backups

Backing up configuration

System > *Maintenance* > *Backup & Restore* enables you to:

- create backup files of the appliance's configuration
- restore the system configuration from a previous backup (see "Restoring a previous configuration" on page 169)
- update the firmware of the FortiDDoS appliance (see "Updating the firmware" on page 81)

Once you have tested your basic installation and verified that it functions correctly, create a backup. This "clean" backup can be used to:

- troubleshoot a non-functional configuration by comparing it with this functional baseline (via a tool such as diff)
- rapidly restore your installation to a simple yet working point (see "Restoring a previous configuration" on page 169)
- batch-configure FortiDDoS appliances by editing the file in a plain text editor, and then uploading the finalized configuration to multiple appliances (see "Restoring a previous configuration" on page 169)

After you have a working deployment, back up the configuration again after any changes. This ensures that you can rapidly restore your configuration exactly to its previous state if a change does not work as planned.

Configuration backups do ***not*** include data such as logs and reports.

For information on backing up DDoS attack log event messages, see "Backing up the DDoS attack log" on page 229.

There are multiple methods that you can use to create a FortiDDoS configuration backup. Use the one that suits your needs:

- "To back up the configuration via the web UI"
- "To back up the configuration via the CLI to a TFTP server"

To back up the configuration via the web UI

1. Log in to the web UI as the `admin` administrator.

2. Go to *System* > *Maintenance* > *Backup & Restore*.

   To access this part of the web UI, your administrator's account access profile must have *Read* and *Write* permission to items in the *Maintenance* category. For details, see "Permissions" on page 44.

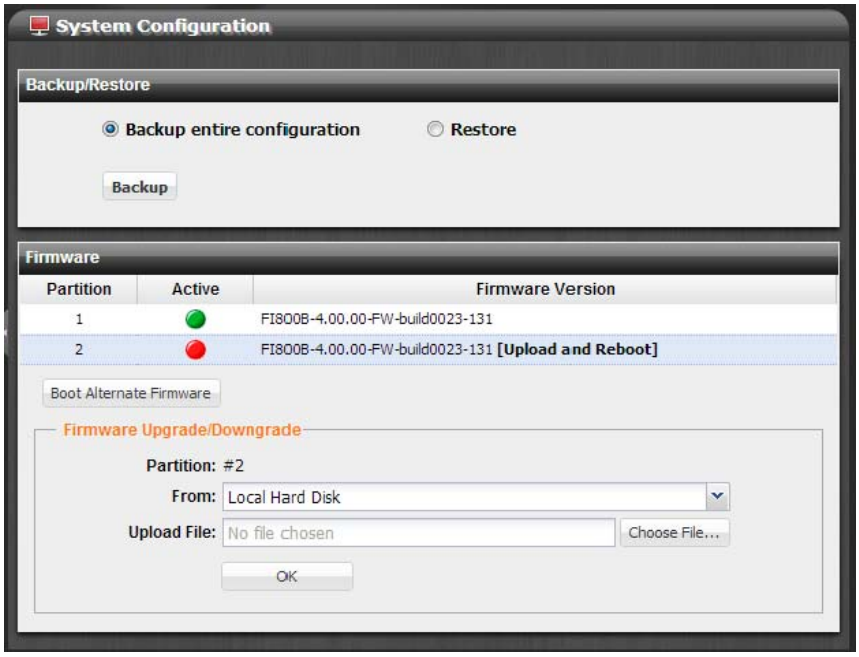

3. Select *Backup entire configuration*, and then click *Backup*.

   If your browser prompts you, navigate to the folder where you want to save the configuration file. Click *Save*.

   Your browser downloads the configuration file. Time required varies by the size of the configuration and the specifications of the appliance's hardware as well as the speed of your network connection, but could take several minutes.

To back up the configuration via the CLI to a TFTP server

1. If necessary, start your TFTP server. (If you do not have one, you can temporarily install and run one such as `tftpd` (Windows, Mac OS X, or Linux) on your management computer.)

   Because TFTP is ***not*** secure, and because it does not support authentication and could allow anyone to have read and write access, you should ***only*** run it on trusted administrator-only networks, ***never*** on computers directly connected to the Internet. If possible, immediately turn off `tftpd` off when you are done.

2. Log in to the CLI as the `admin` administrator using either the local console, the *CLI Console* widget in the web UI, or an SSH or Telnet connection.

   Other administrator accounts do not have the required permissions.

3. Enter the following command:

   ```
   execute backup config tftp <file-name_for_tftp> <tftp_server_ip>
   ```

   where:

Variable	Description
`<file-name_for_tftp>`	Specifies the file name of the backup.
`<server_ipv4>`	Specifies the IP address of the server. **Note:** Domain names are currently ***not*** valid input with this command if you choose the FTP protocol.

   For example, the following command backs up a FortiDDoS 400B's configuration file to a file named `FortiDDoS-400b.conf` in the current directory on the TFTP server 192.0.2.1:

   ```
   FortiDDoS-400b # exec backup config FortiDDoS-400b.conf tftp
   192.0.2.1
   ```

   Time required varies by the size of the database and the specifications of the appliance's hardware, but could take several minutes.

See also

- Restoring a previous configuration

Restoring a previous configuration

If you have backup up your configuration by downloading a configuration file, you can upload the file to revert the appliance's configuration to that point.

Uploading configuration files is also useful for configuring multiple features of the FortiDDoS appliance in a single batch: download a configuration file backup, edit the file in a plain text editor, and then upload the finalized configuration.

To upload a configuration via the web UI

1. Go to *System* > *Maintenance* > *Backup & Restore.*

   To access this part of the web UI, your administrator's account access profile must have *Read* and *Write* permission to items in the *Maintenance* category. For details, see "Permissions" on page 44.

If you have made a configuration backup to an FTP server (see "To back up the configuration via the CLI to a TFTP server" on page 168), you cannot restore it here. Instead, restore it by using the `execute restore` command.

2. Select *Restore*.
3. Do one of the following:
   - For *From File*, specify the path and file name of the file to restore.
   - Click *Browse* to locate the file. (It has a `.conf` file extension.)

4. Click *Restore* to start the restoration of the selected configuration to a file.

   Your web browser uploads the configuration file. Time required to restore varies by the size of the file and the speed of your network connection. The web UI session displays the dashboard (default) content pane.

5. To continue using the web UI, if you have not changed the IP address and static routes of the web UI, simply refresh the web page and log in again.

   Otherwise, to access the web UI again, in your web browser, modify the URL to match the new IP address of the network interface.

   For example, if you configured `mgmt1` with the IP address 192.0.2.1, you would browse to:

   `https://192.0.2.1`

   If the new IP address is on a different subnet than the previous IP address, and your computer is not directly connected to the FortiDDoS appliance, you may also need to modify the IP address and subnet of your computer to match the FortiDDoS appliance's new IP address.

To upload a configuration via the CLI

Use the `execute restore config` command to upload the configuration. For example:

```
execute restore config tftp 440b.conf 172.30.84.105
```

See also

- Backing up configuration

Administrators

In its factory default configuration, FortiDDoS has one administrator account named `admin`. This administrator has permissions that grant full access to FortiDDoS's features.

To prevent accidental changes to the configuration, it is best if only network administrators — and if possible, only a single person — use the `admin` account. You can use the `admin` account to configure more accounts for other people. You can associate each of these accounts with either all Service Protection Profiles or a single protection profile and specify the type of profile settings that each account can access. If you require such role-based access control (RBAC) restrictions, or if you simply want to harden security or prevent inadvertent changes to other administrators' areas, you can do so via access profiles. See "Restricting permissions" on page 174.

For example, you can create an administrator for a specific Service Protection Profile who can make changes to the thresholds and other profile settings. Then create an additional administrator account for the profile that can view the configuration and logs, but ***not*** change them.

Administrators may be able to access the web UI, the CLI, and use ping/traceroute through the network, depending on:

- the account's trusted hosts ("Trusted hosts" on page 46)
- the protocols enabled for each of the FortiDDoS appliance's network interfaces ("Configuring the network interfaces" on page 98)

To configure an administrator account

1. Before configuring the account, configure the access profile that governs the account's permissions (see "Restricting permissions" on page 174).
2. Go to *System* > *Admin* > *Administrators*.

   To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *System* category. For details, see "Permissions" on page 44.
3. Click *Add*.

   A dialog appears.

4. Configure these settings:

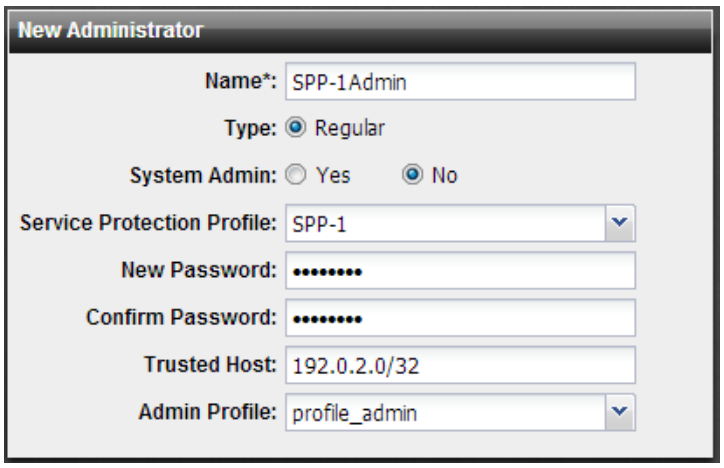

Setting name	Description
Name	Specify the name of the administrator account, such as `admin1`. This is the user name that the administrator provides when logging in to the CLI or web UI. Do not use spaces or special characters, including the 'at' symbol (`@`). The maximum length is 35 characters.
System Admin	Specify whether this administrator is an administrator for all Service Protection Profiles or a single profile.
Service Protection Profile	If this administrator is not a system administrator, select the profile that this account manages.
New Password	Type a password for the administrator account. **Tip:** The password can be any length but cannot be blank. However, for improved security, the password should be at least six characters long, be sufficiently complex, and be changed regularly. To check the strength of your password, you can use a utility such as Microsoft's password strength meter.
Confirm Password	Re-enter the password to confirm its spelling.

Setting name	Description

Trusted Host	Type a source IP address or netmask from which the administrator is allowed to log in to the FortiDDoS appliance. If *PING* is enabled, this is also a source IP address to which FortiDDoS responds when it receives a ping or traceroute signal.
	
	For more information, see "Trusted hosts" on page 46.
	
	To allow login attempts from any IP address, enter `0.0.0.0/0.0.0.0`.
	
	To allow logins only from a single computer, enter its IP address and a 32- or 128-bit netmask:
	
	`192.0.2.2/32`
	
	`2001:0db8:85a3:::8a2e:0370:7334/128`
	
	Caution: If you configure a trusted host, do so for ***all*** administrator accounts. Failure to do so means that all accounts are still exposed to the risk of brute force login attacks. This is because if you leave even ***one*** administrator account unrestricted (that is, `0.0.0.0/0`), the FortiDDoS appliance must allow login attempts on all network interfaces where remote administrative protocols are enabled, and wait until ***after*** a login attempt has been received in order to check the trusted hosts list for that user name.
	
	Tip: If you allow login from the Internet, set a longer and more complex *New Password*, and enable only secure administrative access protocols (*HTTPS* and *SSH*) to minimize the security risk. For information on administrative access protocols, see "Configuring the network interfaces" on page 98.
	
	Tip: For improved security, restrict all trusted host addresses to single IP addresses of computer(s) from which only this administrator logs in.

Admin Profile	Select the access profile that specifies the permissions for this administrator account. For more information on access profiles, see "Restricting permissions" on page 174. You can select *super_admin_prof*, a special access profile used by the `admin` administrator account. However, selecting this access profile does ***not*** confer all permissions of the *admin* administrator. For example, the new administrator cannot reset lost administrator passwords. This option does not appear for the `admin` administrator account, which by definition always uses the *super_admin_prof* access profile.

5. Click *Save*.

See also

- Restricting permissions
- Configuring the network interfaces
- Trusted hosts
- Permissions

Restricting permissions

Access profiles determine administrator accounts' permissions.

Administrators who have only read access to a feature can perform the following tasks:

- Access the web UI page for that feature
- Use the `get` and `show` CLI command for that feature

Read-only administrators cannot perform the following tasks:

- Make changes to the configuration. When they click *Add*, *Save*, *Edit* or other modification commands, FortiDDoS displays a permission denied message. Write access is required for modification of any kind.
- Access `config` CLI commands. Only `get` or `show` commands are available to read-only administrators.

In larger companies where multiple administrators share the work, access profiles often reflect the specific job that each administrator does ("role"), such as account creation or log auditing. Access profiles can limit each administrator account to their assigned role. This is sometimes called role-based access control (RBAC).

The `super_admin_prof` access profile, a special access profile assigned to the `admin` administrator account and required by it, ***does not*** appear in the list of access profiles. It exists by default and cannot be changed or deleted, and consists of essentially UNIX `root`-like permissions.

Even if you assign the `super_admin_prof` access profile to other administrators, they will ***not*** have all of the same permissions as the `admin` account. The `admin` account has some special permissions, such as the ability to reset administrator passwords, that are inherent in that account only. Other accounts should not be considered a complete substitute.

If you create more administrator accounts, whether to harden security or simply to prevent accidental modification, create other access profiles with the minimal degrees and areas of access that each role requires. Then assign each administrator account the appropriate role-based access profile.

For example, for an administrator whose only role is to audit the log messages, you might make an access profile named `auditor` that only has *Read* permissions for the *Log & Report* area.

To configure an access profile

1. Go to *System* > *Admin* > *Access Profile*.

   To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *System* category. For details, see "Permissions" on page 44.
2. Click *Add*.

   A dialog appears.
3. In *Profile Name*, type a unique name that can be referenced by other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters.

4. Configure the permissions options.

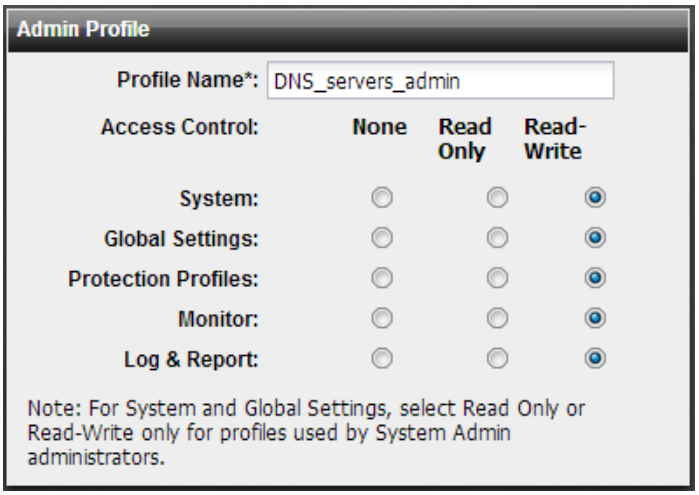

   For each row associated with an area of the configuration, select either the *Read Only* or *Read-Write* radio buttons to grant that type of permission. For a list of features governed by each access control area, see "Permissions" on page 44.

5. Click *Save.*

See also

- Administrators
- Permissions

Changing an administrator's password

If an administrator has forgotten or lost his or her password, or if you need to change an administrator account's password and you do not know the current password, you can reset the password.

If you forget the password of the `admin` administrator, you can reset the FortiDDoS appliance to its default state (including the default administrator account and password) by restoring the firmware. For instructions, see "Restoring firmware ("clean install")" on page 267.

To change an administrator account's password

1. Log in as the `admin` administrator account.

   Alternatively, if you know the current password for the account whose password you want to change, you may log in with any administrator account whose access profile permits *Read-Write* access to items in the *System* category.

2. Go to *System* > *Admin* > *Administrators*.
3. Double-click the row of the account whose password you want to change.

4. Mark the check box named *Change Password*.

   Additional text fields appear where you can enter the new password.

5. In the *Old Password* field, if available, type the current password for the account. (The `admin` account does not have an old password initially.)

   This field does not appear for other administrator accounts if you are logged in as the `admin` administrator.

6. In the *New Password* and *Confirm Password* fields, type the new password. It should have sufficient complexity and number of characters to deter brute force and other attacks.

   The web UI verifies that the passwords match exactly to ensure that there have been no typos, and alerts you with a red icon to the right of the field if the passwords do no match. You cannot save the settings if the passwords do not match.

7. Click *Save*.

   The new password takes effect the next time that administrator account logs in.

Service Protection Profile settings

Creating an SPP is part of setting up your FortiDDoS. For detailed instructions, see “Create a service protection profile (SPP)” on page 113.

Click *Protection Profiles* > *SPP Settings* > *SPP Settings* to configure the general operating mode as well as specific threat mitigation features for individual Service Protection Profiles (SPPs).

Setting name	See
Inbound operating mode **Outbound operating mode**	“Setting FortiDDoS to detection mode” on page 116
SYN flood mitigation direction **SYN flood mitigation mode**	“Configuring SYN flood mitigation feature controls” on page 136
Adaptive mode **Adaptive limit (in percentage)**	“See also” on page 138
Source penalty factor inbound/outbound **Application penalty factor inbound/outbound** **Destination penalty factor inbound/outbound**	“Setting penalty factors” on page 177
Mandatory http header count	“Set the mandatory HTTP header count” on page 179
TCP session feature control	“Configuring TCP session feature control” on page 179
Aggressive aging TCP connections feature control	“Configuring aggressive aging feature controls” on page 181

Setting penalty factors

Penalty factors allow you to adjust how FortiDDoS tracks traffic that is associated with specific types of flood attacks. When you configure penalty factors, FortiDDoS multiplies the packet volume for specific types of traffic by the factor value. As a result, this type of traffic reaches the specified threshold sooner than other types of traffic. You can use these factors to increase the appliance’s sensitivity to the following types of traffic:

- Traffic from a source that FortiDDoS has identified as the source of a flood
- HTTP (Layer 7 floods associated with a URL or Host, Referer, Cookie or User-Agent header field)
- Traffic to a destination that is under attack

Because incoming traffic is more likely to be the source of a threat, you can configure different penalties for incoming and outgoing traffic.

In addition to the source penalty factor, you can apply an application penalty factor. For example, if *Source penalty factor* is 2 and *Application penalty factor* is 8, the total penalty used to track HTTP (Layer 7) traffic is 2*8, or 16.

For example, in the case of a User Agent flood attack, a source is sending a User Agent that is overloaded. Each time the source sends a layer 7 packet with the User Agent, FortiDDoS considers the packet equivalent to 16 packets.

The destination penalty factor helps you to identify destinations that may be under attack by increasing the count of packets for a specific destination. When you specify this factor, incoming or outgoing traffic reaches the threshold that regulates traffic to individual destinations more quickly.

When you specify a destination penalty factor that is too low, FortiDDoS does not react quickly enough to detect and block traffic that is attacking a destination. If the value is too high, the appliance may limit or block traffic to destinations that are not under attack.

To set penalty factors via the web UI

1. Click *Protection Profiles* > *SPP Settings* > *SPP Settings*, and then configure these settings:

Setting name	Description
Source penalty factor inbound/outbound	Applies the specified penalty to traffic to and from any source of a flood attack. Default is 1 (no penalty).
Application penalty factor inbound/outbound	Applies the specified penalty to traffic to and from sources associated with HTTP in addition to the source penalty factor. HTTP includes elements such as a URL, Host, Referer, Cookie or User Agent. Default is 1 (no penalty)
Destination penalty factor inbound/outbound	Applies the specified penalty to traffic for an individual destination. Default is 1 (no penalty)

2. Click *Save*.

To set penalty factors via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp setting
    set source-penalty-factor-inbound <integer>
    set source-penalty-factor-outbound <integer>
    set application-penalty-factor-inbound <integer>
    set application-penalty-factor-outbound <integer>
    set destination-penalty-factor-inbound <integer>
    set destination-penalty-factor-outbound <integer>
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `<integer>` specifies the factor by which the specified type of traffic is increased when FortiDDoS is evaluating traffic for attacks

See also

- Adjusting thresholds

Set the mandatory HTTP header count

Mandatory http header count specifies the minimum number of HTTP header fields that are required in a HTTP GET request for a URL.

You can also specify the maximum number of packets that do not have the mandatory HTTP header fields a single source can send. For information on this threshold and HTTP header fields, see "Adjusting thresholds individually" on page 155.

To set the mandatory HTTP header count via the web UI

1. Go to *Protection Profiles* > *SPP Settings* > *SPP Settings*.
2. For *Mandatory http header count*, enter the minimum number of header fields that FortiDDoS requires a GET request to have.
3. Click *Save*.

To set the mandatory HTTP header count via the CLI

Enter the following commands:

```
edit <spp_name>
   config ddos spp setting
      set mandatory-http-header-count <integer>
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `<integer>` specifies the minimum number of header fields that FortiDDoS requires a GET request to have

See also

- Adjusting thresholds individually

Configuring TCP session feature control

TCP is a stateful protocol. FortiDDoS allows you to allow or deny the following anomalies related to TCP states:

- Sequence Validation

  The TCP state machine ensures that TCP sequence numbers for the packets within a session are valid.
- SYN Validation

  The TCP state machine allows only TCP SYNs from IP addresses in the legitimate IP addresses (for example, sources have done a 3-way handshake in the past).
- Track State Transition Anomalies

  The TCP state machine ensures that TCP state transitions follow the rules. For example, if an ACK packet is received when FortiDDoS has not observed a SYN/ACK packet, it is a state transition anomaly.

- Foreign Packet Validation

  The TCP state machine drops TCP packets without an existing TCP connection and reports them as a foreign packet.

  Here are some reasons why foreign packets occur when they shouldn't:

  - Detection mode

    When the appliance is in detection mode, it can forward packets or sessions but think that it dropped them. When subsequent packets arrive for that connection, the appliance treats them as foreign packets because the associated session is closed.

  - Timeout differences between servers and the FortiDDoS appliance

    The appliance TCP session timeout is lower than the one for typical Windows or Linux servers. This configuration allows FortiDDoS to handle more new sessions per second. But sometimes this difference creates time differences and FortiDDoS drops packets that appear to come from sessions that FortiDDoS has purged. Most of the time, these packets are not data packets. They may be simple terminating packets.

  - HTTP Browser Behavior

    Most people navigate from site to site without closing the browser. The browser sends RST or FIN packets to close the connections to the site only when it is closed. The appliance sessions may have been purged by the time RST or FIN packets arrive and there may not be any data associated with this session - just an RST packet.

  In most cases, the foreign packets validation is useful for filtering out junk and enabling it is not important. Because the number of foreign packets is pretty high, FortiDDoS does not store the source and destination of each packet. Therefore, you may not be able to determine the origin of a foreign packet.

- Allow Tuple Reuse

  The TCP state machine updates the TCP entry when a tuple is reused. This update occurs only during the closed or close-wait, fin-wait, time-wait states, when the connection is just about to retire.

- Allow Duplicate SYN in SYN_SENT

  The TCP state machine allows duplicate TCP SYN packets during the SYN-SENT state. It allows this type of packet even if the sequence numbers are different.

- Allow Duplicate SYN in SYN_RECV

  The TCP state machine allows duplicate TCP SYN packets during the SYN-RECV state. It allows this type of packet even if the sequence numbers are different.

- Allow SYN Anomaly, Allow SYN/ACK Anomaly, Allow ACK Anomaly, Allow RST Anomaly, Allow FIN Anomaly

  The TCP state machine allows duplicate TCP packets during any other state even if the sequence numbers are different from the existing connection entry. This is equivalent to allowing the packet without updating an existing connection entry with the new information from the allowed packet.

To select which TCP state anomalies FortiDDoS allows or denies via the web UI

1. Click *Protection Profiles* > *SPP Settings* > *SPP Settings*.
2. For *TCP session feature control*, click the down arrow beside the field to display the list of options.
3. Click an option to select it. Click a selected option to clear it.
4. Click *Save*.

To select which TCP state anomalies FortiDDoS allows or denies via the CLI

Enter the following commands:

```
edit <spp_name>
config ddos spp setting
     set tcp-session-feature-control
        {sequence-validation
        syn-validation
        state-transition-anomalies-validation
        foreign-packet-validation
        allow-tuple-reuse
        allow-duplicate-syn-in-syn-sent
        allow-duplicate-syn-in-syn-recv
        allow-syn-anomaly
        allow-syn-ack-anomaly
        allow-ack-anomaly
        allow-rst-anomalyallow-fin-anomaly}
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `{sequence-validation syn-validation...}` specifies the options that are selected

Configuring aggressive aging feature controls

Aggressive aging helps to protect servers from connection overload. When some servers are under attack from numerous slow connections, they get overloaded. You can relieve them by aggressively aging the connections based on specific criteria, such as aging connections that are sitting idle.

The following aggressive aging options are available:

- url-flood

  When a specific URL gets overloaded (URL flood), FortiDDoS sends an RST (Reset) packet to servers to relieve them of excessive connection overload.

  Enabled by default.

- old-tcp-connections

  FortiDDoS ages out from its internal memory connections that are idle (no data transfer) for the specified length of time. It also relieves the server by sending RST/ACK to it.

  The default timeout is 60 seconds. The maximum value allowed is 36 hours (specified in seconds). (See “To change the slow connection byte count threshold and slow connection timeout” on page 183.)

- slow-data-transfer

  FortiDDoS ages out any connection that sends an amount of traffic that is smaller than the value of the slow connection byte count threshold during the specified observation period. It ages out these connections by sending a “reset” (RST) flag to the server for the connection.

  By default, the threshold value is 16 * 64 bytes (1024 bytes) and the timeout value is 60 seconds.

  You specify the byte count threshold using units of 64 bytes. A value of 1 means 64 bytes, 2 means 128 bytes, and so on. The maximum value allowed is 160 * 64 bytes (10 240 bytes).

  You specify the slow connection timeout in seconds. The maximum value allowed is 129 600 seconds (36 hours).

  For information on how to change this threshold and timeout, see “To change the slow connection byte count threshold and slow connection timeout”.

  To generate information about the slow connections that FortiDDoS drops that you can view in a report, select source-track-slow-tcp-connections. For more information on tracking these dropped connections, see “Tracking slow data connections that FortiDDoS has aged out”.

  This setting is designed to counter attacks such as slowloris or R.U.D.Y. (The slowloris program attempts to open a large number of server connections and keep those connections open for as long as possible. The R.U.D.Y. tool uses long form field submissions to consume server resources.) For example, you can stop an attack that uses the default slowloris script by setting the slow connection byte count threshold to the recommeded value of 20 (that is, 1280 bytes) and the slow connection timeout to 10 or 15 seconds (also the recommended values).

- high-concurrent-connection-per-source, high-concurrent-connection-per-destination

  FortiDDoS ages out any connection that is idle (that is, one that transfers no data) during the time period specified by the idle connections timeout. FortiDDoS ages out these connections by sending RST/ACK to the server for the connection.

  The default timeout value is 60 seconds. The maximum value allowed is 129 600 seconds (36 hours). (See “To change the slow connection byte count threshold and slow connection timeout” on page 183.)

  Enabled by default.

- source-track-slow-tcp-connections

  When this option and slow-data-transfer are enabled, FortiDDoS generates report information about the slow connections that it drops.

For more information, see “Tracking slow data connections that FortiDDoS has aged out”.

To configure the aggressive aging features via the web UI

1. Click *Protection Profiles* > *SPP Settings* > *SPP Settings*.
2. For *Aggressive aging TCP connections feature control*, click the down arrow beside the field to display the options.
3. Click an option to select it. Click a selected option to clear it.
4. Click *Save*.

To configure the aggressive aging features via the CLI

Enter the following commands:

```
edit <spp_name>
config ddos spp setting
   set aggressive-aging-feature-control
   {url-flood
```

```
    old-tcp-connections
    slow-data-transfer
    high-concurrent-connection-per-source
    high-concurrent-connection-per-destination
    source-track-slow-tcp-connections}
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `{url-flood old-tcp-connections...}` specifies the options that are selected

To change the slow connection byte count threshold and slow connection timeout

Do one of the following:

- Click *Global Settings* > *Settings* > *Settings*, and then enter new values for the following settings:
  - *Slow TCP connection byte threshold (multiples of 64-byte)*
  - *Slow TCP connection observation period (in seconds)*
- Enter the appropriate commands using the command line interface (CLI). For example, to set the threshold to 128 bytes and the timeout to 2 minutes, enter:

```
config ddos global setting
  set slow-tcp-connection-byte-threshold 2
  set slow-tcp-connection-observation-period 120
end
```

Tracking slow data connections that FortiDDoS has aged out

When you enable the *slow-data-transfer* option, FortiDDoS drops packets from slow data connections. FortiDDoS does not generate report information about these dropped packets except in the following situations:

- The Foreign Packet Validation option is enabled.

  When FortiDDoS detects a packet for a connection that it has closed, it records a foreign packet drop event in the *DDoS Attack Log* (State anomaly event).
- The *source-track-slow-tcp-connections* option is enabled.

  FortiDDoS blocks the IP address of the source of the slow connections and records the event in the *DDoS Attack Log* (Source flood).

MAC address for aggressive aging

By default, FortiDDoS uses the MAC address 00:00:00:00:00:00 for its aggressive aging features. However, some firewalls block this address. To configure a different MAC address, go to *Global Settings* > *Settings* > *Settings*.

Advanced/optional system settings

The *System* menu configures a variety of settings that apply to the entire FortiDDoS appliance.

Many system settings must be configured during the initial installation.For required system settings, see the appropriate section of “How to set up your FortiDDoS” on page 53.

Changing the FortiDDoS appliance’s host name

The host name of the FortiDDoS appliance is used in several places.

- The name appears in the *System Information* widget on *System* > *Status* > *Dashboard*. For more information about the *System Information* widget, see “System Information widget” on page 196.
- It is used in the command prompt of the CLI.
- It is used as the SNMP system name. For information about SNMP, see “SNMP traps & queries” on page 243.

The *System Information* widget and the `get system status` CLI command display the full host name. If the host name is longer than 16 characters, FortiDDoS may display a tilde (~) instead of the additional characters.

For example, if the host name is FortiDDoS123456789, the CLI prompt is `FortiDDoS1234567~#`.

Administrators whose access profiles permit *Write* access to items in the *System* category can change the host name.

To change the host name of the FortiDDoS appliance

1. Go to *System* > *Status* > *Dashboard*.
2. In the *System Information* widget, in the *Host Name* row, click *Change*.
3. In the *New Name* field, type a new host name.

   The host name can be up to 35 characters in length. It can include US-ASCII letters, numbers, hyphens, and underscores, but ***not*** spaces and special characters.
4. Click *Save*.

See also

- System Information widget

Global Settings dialog box

Click *Global Settings* > *Settings* > *Settings* to access settings that configure system-wide functionality or apply to all Service Protection Profiles (SPPs).

Link down synchronization	See “Configuring bypass mode” on page 186.
Slow TCP connection byte threshold (multiples of 64-byte) **Slow TCP connection observation period (in seconds)**	See “Configuring aggressive aging feature controls” on page 181.
Blocking Period for all attacks (in seconds) **Blocking Period for Identified Sources (in seconds)** **Extended blocking Period for Identified Sources (in seconds)** **Drop Threshold to extend blocking period for Identified Sources**	See “Configuring blocking periods” on page 137.
HTTP URL per source tracking period(in seconds)	Specifies the observation period that FortiDDoS uses to determine if a source has exceeded the maximum number of HTTP accesses specified by the http-urls-per-source threshold.
HTTP source tracking observation period(in seconds)	Specifies the observation period that FortiDDoS uses to determine if a source has exceeded the maximum number of times a single IP address can retrieve the same URL back-to-back (that is, without accessing other URLs in between) specified by the http-sequential-access-per-source threshold.
Source MAC address for aggressive aging	Specify the MAC address that FortiDDoS uses for its aggressive aging features. By default, FortiDDoS uses the MAC address 00:00:00:00:00:00 for its aggressive aging features. However, some firewalls block this address. For more information, see “Configuring aggressive aging feature controls” on page 181.
HTTP Anomaly	See “Configuring HTTP anomaly features” on page 187.
IPv6 dual stack support	See “Enabling Internet Protocol version 6 (IPv6) support” on page 107.
Geo Location Policy	See “Blocking addresses from a specific geographic location, anonymous proxies, and satellite providers” on page 120.

Configuring bypass mode

The FortiDDoS bypass feature is designed to maintain the flow of traffic through the appliance if its monitoring and protection functionality fails. However, if the data path fails (for example, a power failure), it maintains data traffic on copper links only and not on fiber-optic links. You can use an external bypass switch to maintain connectivity during a data path failure. For more information, see "External bypass switches for maintenance & failover" on page 55

Configuring link down synchronization or link state propagation

FortiDDoS monitors the state of each port pair — the pair of physically linked ports that receive and transmit regulate incoming and outgoing traffic. This pair corresponds to an odd and even-numbered port (for example, Port 1 connected to the Ethernet and Port 2 connected to the Internet).

If the link goes down on either port and *Link Down Synchronization* is set to *Wire* (the default value), FortiDDoS automatically disables the other port in the link pair. The appliance re-enables the port when it detects that the link for other port in the pair is up again.

To disable this synchronization, set *Link Down Synchronization* to *Hub*.

Changing the *Link Down Synchronization* option restarts FortiDDoS. Please plan for this downtime.

To configure link down synchronization via the web UI

1. To configure *Link Down Synchronization*, click *Global Settings* > *Settings* > *Settings*.
2. For *Link Down Synchronization*, select *Wire* or *Hub*.
3. Click *Save*.

To configure link down synchronization via the CLI

Enter the following commands:

```
config ddos global setting
  set link-down-synchronization {wire | hub}
end
```

where:

- `{wire | hub}` specifies whether FortiDDoS automatically disables a port in a port pair when the other port in the pair is down

See also

- External bypass switches for maintenance & failover

Configuring HTTP anomaly features

The *HTTP Anomaly* option allows you to enable one or more features that block or drop HTTP packet traffic based on the HTTP methods it uses:

known-opcode-anomaly	For future use. Currently, FortiDDoS does not allow you to restrict traffic to specific HTTP methods.
unknown-opcode-anomaly	Drops any HTTP traffic that uses a method other than one of the following: GET, HEAD, OPTIONS, PUT, POST, CONNECT, DELETE, or TRACE. For example, TEST or PROPFIND. Generates the attack log message 'Unknown HTTP Anomaly'.
invalid-opcode-anomaly	Drops any HTTP traffic with an HTTP version other than one of the following: 0.9, 1.0, or 1.1. Generates the attack log message 'Invalid HTTP Version Anomaly'.
http-version-0-9	By default, FortiDDoS drops any HTTP traffic that uses HTTP version 0.9. Select this option to allow version 0.9 traffic.

Certificate configuration

When a FortiDDoS appliance initiates or receives an SSL or TLS connection, it uses certificates. FortiDDoS uses certificates for management network interface connections for:

- encryption (for example, HTTPS)
- authentication of the appliance by clients
- authentication of the appliance by FortiGuard servers when downloading of IP reputation lists

In most cases, you do not need to change the certificate because the appliance has a built-in key pair and certificate. This built-in certificate cannot be verified by a trusted CA.

However, you can install a new certificate, if required, by using FortiDDoS to generate a key pair and acquiring a certificate from a signing authority.

See also

- Generating a certificate signing request
- Uploading a certificate
- How to export/back up certificates & private keys

Generating a certificate signing request

FortiDDoS allows you to generate your own certificate signing request (CSR). A CSR is an unsigned certificate file that the CA signs. When FortiDDoS generates the CSR, it also

generates the associated private key that the appliance can use to sign and encrypt connections with clients. You submit the CSR for verification and signing by the CA.

To generate a certificate request

1. Go to *System* > *Certificates* > *Local*.

   To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *System* category. For details, see “Permissions” on page 44.

2. Click *Generate*.

   A dialog appears.

3. Configure the certificate signing request:

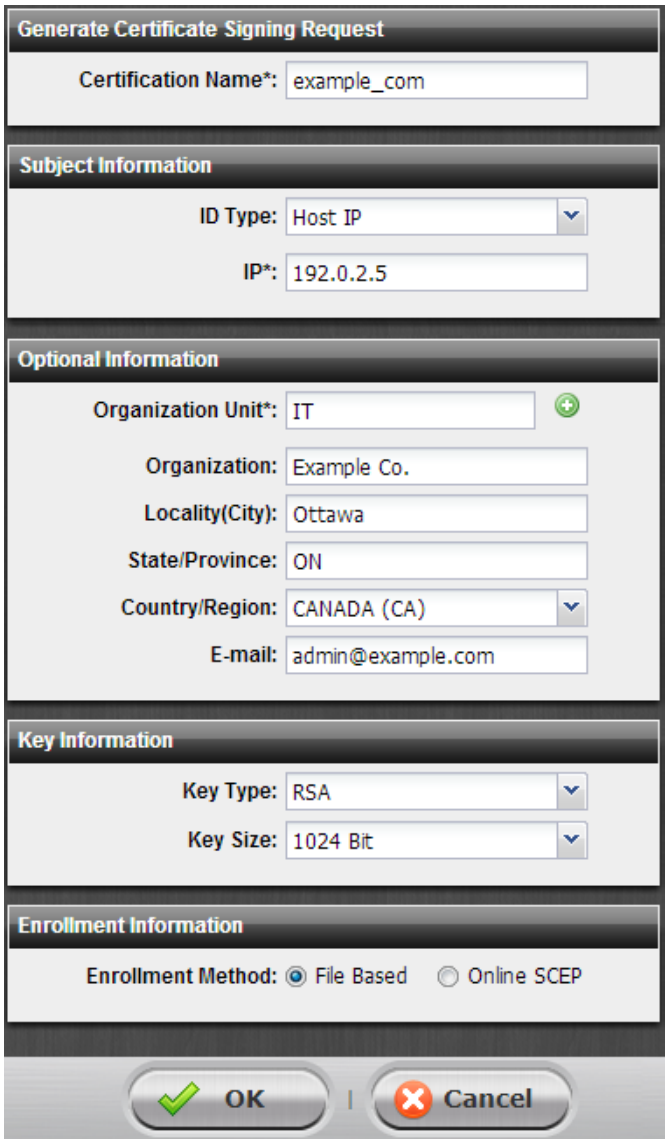

Setting name	Description
Certification Name	Type a unique name for the certificate request file, such as `www_example_com`, that can be referenced in other parts of the configuration. Do not use spaces or special characters. The maximum length is 35 characters. **Note:** This is the name of the CSR file, not the host name/IP contained in the certificate's `Subject:` line.

Setting name	Description
Subject Information	Includes information that the certificate is required to contain in order to uniquely identify the FortiDDoS appliance. This area varies depending on the *ID Type* selection.
ID Type	Select the type of identifier to use in the certificate to identify the FortiDDoS appliance: • **Host IP** — Type the ***static*** public IP address of the FortiDDoS appliance in the *IP* field. If the FortiDDoS appliance does not have a static public IP address, use *E-Mail* or *Domain Name* instead. **Note:** If your network has a dynamic public IP address, do not use this option. Users' browsers display "Unable to verify certificate" or similar error message when your public IP address changes. • **Domain Name** — Type the FQDN of the FortiDDoS appliance, such as `www.example.com`, in the *Domain Name* field. This does not require that the IP address be static, and may be useful if, for example, your network has a dynamic public IP address and therefore clients connect to it via dynamic DNS. Do not include the protocol specification (http://) or any port number or path names. • **E-Mail** — Type the email address of the owner of the FortiDDoS appliance in the *E-mail* field. Use this if the appliance does not require either a static IP address or a domain name. The type you should select varies by whether or not your FortiDDoS appliance has a static IP address, a fully-qualified domain name (FQDN), and by the primary intended use of the certificate. For example, if your FortiDDoS appliance has both a static IP address and a domain name, but you prefer to make HTTPS connections to the web UI by the domain name, you might prefer to generate a certificate based upon the domain name of the FortiDDoS appliance, rather than its IP address. Depending on your choice for *ID Type*, related options appear.
IP	Type the static IP address of the FortiDDoS appliance, such as `10.0.0.1`. The IP address should be the one that is visible to clients. Usually, this should be its public IP address on the Internet, or a virtual IP that you use NAT to map to the appliance's IP address on your private network. This option appears only if *ID Type* is *Host IP*.

Setting name		Description
	Domain Name	Type the fully qualified domain name (FQDN) of the FortiDDoS appliance, such as `www.example.com`. The domain name must resolve to the IP address of the FortiDDoS appliance according to the DNS server used by clients. (If it does not, the clients' browsers display a *Host name mismatch* or similar error message.) For more information, see "Configuring the network interfaces" on page 98. This option appears only if *ID Type* is *Domain Name*.
	E-mail	Type the email address of the owner of the FortiDDoS appliance, such as `admin@example.com`. This option appears only if *ID Type* is *E-Mail*.
Optional Information		Includes information that you may include in the certificate, but which is not required.
	Organization Unit	Type the name of your organizational unit (OU), such as the name of your department. This is optional. To enter more than one OU name, click the + icon, and enter each OU separately in each field.
	Organization	Type the legal name of your organization. This is optional.
	Locality (City)	Type the name of the city or town where the FortiDDoS appliance is located. This is optional.
	State/Province	Type the name of the state or province where the FortiDDoS appliance is located. This is optional.
	Country/Region	Select the name of the country where the FortiDDoS appliance is located. This is optional.
	E-mail	Type an email address that may be used for contact purposes, such as `admin@example.com`. This is optional.
Key Type		Displays the type of algorithm used to generate the key. This option cannot be changed, but appears in order to indicate that only RSA is currently supported.
Key Size		Select a secure key size of *512 Bit*, *1024 Bit*, *1536 Bit* or *2048 Bit*. Larger keys are slower to generate, and make FortiDDoS slower to decrypt/encrypt each packet that uses the key, but provide better security.
Enrollment Method		Select either: • **File Based** — Requires you to manually download and submit the resulting certificate request file to a certificate authority (CA) for signing. Once signed, upload the local certificate. • **Online SCEP** — The FortiDDoS appliance automatically uses HTTP to submit the request to the simple certificate enrollment protocol (SCEP) server of a CA, which validates and signs the certificate. For this selection, two options appear. Enter the *CA Server URL* and the *Challenge Password*.

4. Click *Save*.

   The FortiDDoS appliance creates a private and public key pair. The generated request includes the public key of the FortiDDoS appliance and information such as the FortiDDoS appliance's IP address, domain name, or email address. The FortiDDoS appliance's private key remains confidential on the FortiDDoS appliance. The *Status* column of the new CSR entry is *Pending*.
5. Select the row that corresponds to the certificate request.
6. Click *Download*.

   Standard dialogs appear with buttons to save the file at a location you select. Your web browser downloads the certificate request (.csr) file. Time required varies by the size of the file and the speed of your network connection.
7. Upload the certificate request to your CA.

   After you submit the request to a CA, the CA verifies the information in the certificate, gives it a serial number and an expiration date, and signs it with the public key of the CA.
8. If you are not using a commercial CA whose root certificate is already installed by default on web browsers, download your CA's root certificate, and then install it on all computers that connect to your appliance. (If you do not install these, those computers may not trust your new certificate.)
9. When you receive the signed certificate from the CA, upload the certificate to the FortiDDoS appliance (see "Uploading a certificate" on page 192).

See also

- Uploading a certificate

Uploading a certificate

You can import (upload) either:

- Base64-encoded
- PKCS #12 RSA-encrypted

X.509 certificates and private keys to the FortiDDoS appliance.

If a certificate is signed by an intermediate certificate authority (CA) rather than a root CA, before clients trust the certificate, you must demonstrate a link with root CAs that the clients trust, thereby proving that the certificate is genuine. You can demonstrate this chain of trust by appending a signing chain in the certificate.

To append a signing chain in the certificate itself, before uploading the server certificate to the FortiDDoS appliance

1. Open the certificate file in a plain text editor.

2. Append the certificate of each intermediary CA in order from the intermediary CA who signed the local certificate to the intermediary CA whose certificate was signed directly by a trusted root CA.

   For example, a certificate that includes a signing chain might use the following structure:

```
-----BEGIN CERTIFICATE-----
<server certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<certificate of intermediate CA 1, who signed the server certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<certificate of intermediate CA 2, who signed the certificate of
intermediate CA 1 and whose certificate was signed by a trusted root
CA>
-----END CERTIFICATE-----
```

3. Save the certificate.

To upload a certificate

1. Go to *System* > *Certificates* > *Local*.

   To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *System* category. For details, see "Permissions" on page 44.

2. Click *Import*.

   A dialog appears.

3. Configure these settings:

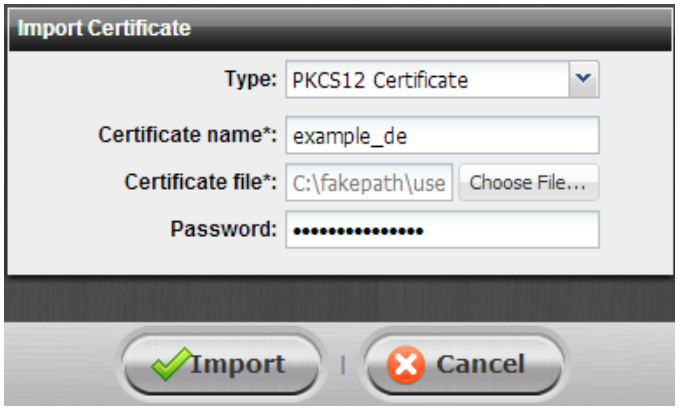

Setting name	Description
Type	Select the type of certificate file to upload, either: • **Local Certificate** — An unencrypted certificate in PEM format. • **Certificate** — An unencrypted certificate in PEM format. The key is in a separate file. • **PKCS12 Certificate** — A PKCS #12 password-encrypted certificate with key in the same file. Other fields may appear depending on your selection.
Certificate name	Type a name that can be referenced by other parts of the configuration, such as `www_example_com`. Do not use spaces or special characters. The maximum length is 35 characters.

Certificate file	Click *Browse* to locate the certificate file that you want to upload. This option is available only if *Type* is *Certificate* or *Local Certificate*.
Key file	Click *Browse* to locate the key file that you want to upload with the certificate. This option is available only if *Type* is *Certificate*.
Certificate with key file	Click *Browse* to locate the PKCS #12 certificate-with-key file that you want to upload. This option is available only if *Type* is *PKCS12 Certificate*.
Password	Type the password that was used to encrypt the file. The FortiDDoS appliance will use the password to decrypt and install the certificate. This option is available only if *Type* is *Certificate* or *PKCS12 Certificate*.

4. Click *OK*.

See also

- Generating a certificate signing request

How to export/back up certificates & private keys

Because your X.509 certificates and private keys are vital to the FortiDDoS configuration, they are included in each FortiDDoS backup (see “Backing up configuration” on page 167). Should FortiDDoS experience hardware failure, this backup minimizes time required for you to reconfigure a replacement appliance.

Monitoring attack activity and other system information

To get the most value out of your FortiDDoS appliance, use it to keep informed about your network, not just to protect it. FortiDDoS appliances have many tools that you can use to monitor the status of systems, attack activity, uptime, and throughput.

Traffic statistics reports allow you to view the maximum packet rate for each traffic parameter for a specific Service Protection Profile (SPP) and time period, see "Generating and reviewing a traffic statistics report" on page 144.

See also

- The dashboard
- Traffic graphs
- Logging
- Alert email
- SNMP traps & queries
- Reports
- Attack Graphs dashboard

The dashboard

System > *Status* > *Dashboard* is displayed when you log in to the web UI. The dashboard widgets indicate performance level, attack activity, or the status of other system components.

Each day, check the dashboard for obvious problems.

By default, the dashboard contains the following widgets:

- System Information widget
- License Information widget
- CLI Console widget
- SPP Attacks widget
- Event Log Console widget
- System Status widget
- Count of Unique Sources widget
- System Resources widget

Figure 26: Viewing the dashboard (*System* > *Status* > *Dashboard*)

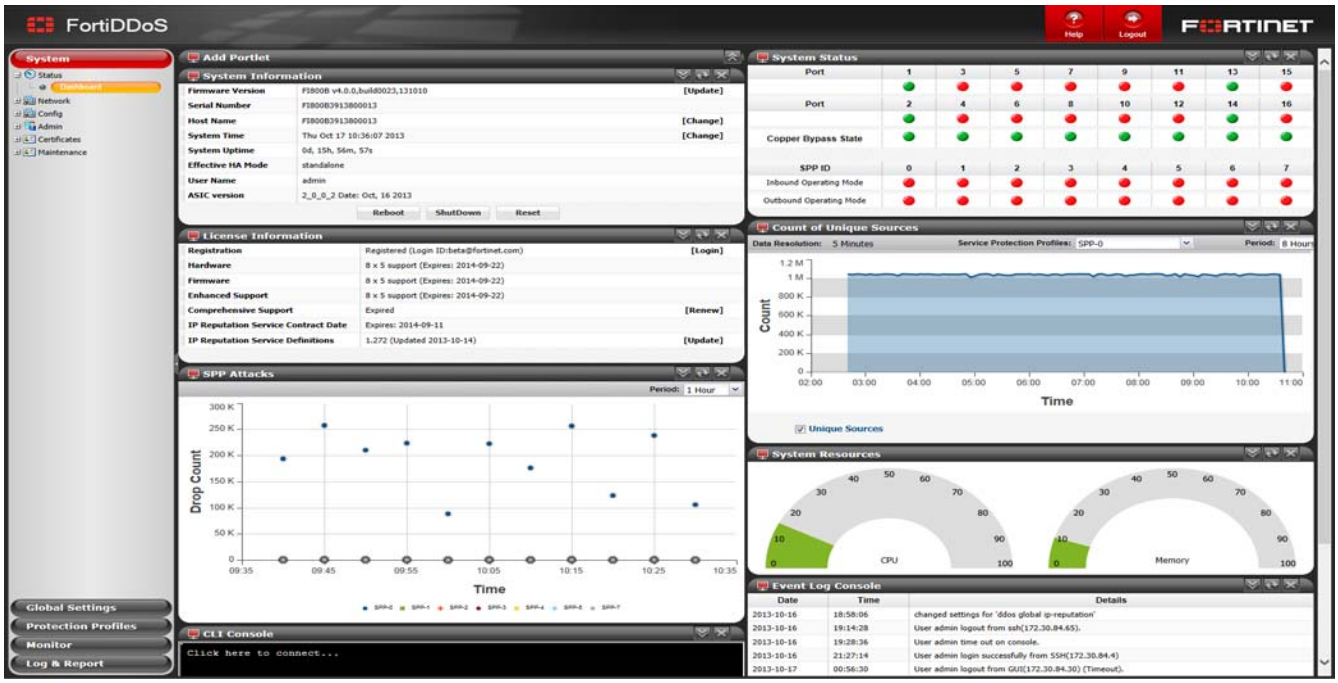

In the default dashboard setup, widgets display the current system status of the FortiDDoS appliance, including host name, firmware version, system time, port status, and current counts of dropped and blocked packets. The dashboard also contains a CLI widget that enables you to use the command line interface (CLI) through the web UI.

You can customize the dashboard by selecting which widgets to display, where they are located on the page, and whether they are minimized or maximized.

Drag the widget's title bar to move it to a new location.

To display any of the widgets not currently shown on *System* > *Status* > *Dashboard*, expand *Add Portlet* widget and select the widget you want to add. Because FortiDDoS can display only one instance of a widget on the dashboard, any widgets that are already displayed are not available in the *Add Portlet* menu.

The widget's title bar contains buttons to close, minimize or maximize the widget.

Figure 27:A minimized widget

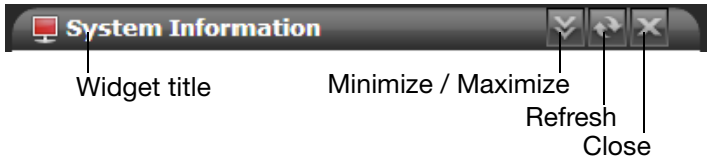

To access the dashboard, your account access profile must have *Read* permission to items in the *System* category. To use features that alter the FortiDDoS or perform actions, you may also need *Write* permissions in various categories. For details, see "Restricting permissions" on page 174.

System Information widget

The *System Information* widget on the dashboard displays basic information, such as the firmware version, host name, and system time.

It also enables you to configure some basic attributes such as the host name and to update the firmware.

FortiDDoS administrators whose access profiles permit *Write* access to items in the *System* category can change the firmware, host name, and system time.

Table 10: *System Information* widget

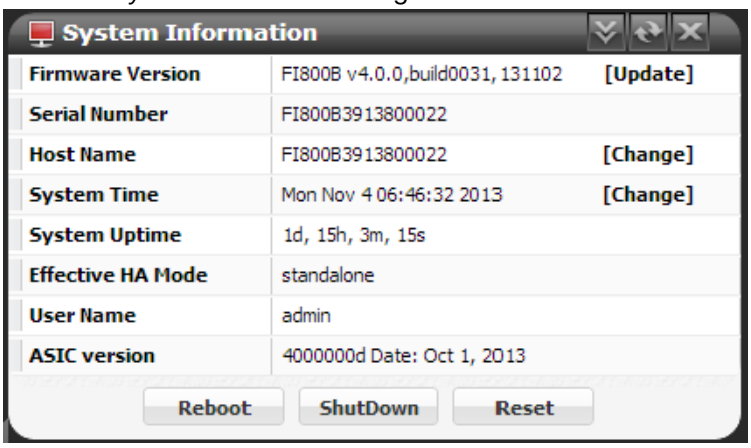

Setting name	Description
Firmware Version	Displays the version of the firmware currently installed on the FortiDDoS appliance. Click *Update* to install a new version of firmware. See "Updating the firmware" on page 81.
Host Name	Displays the host name of the FortiDDoS appliance. Click *Change* to change the host name. See "Changing the FortiDDoS appliance's host name" on page 184.
System Time	Displays the current date and time according to the FortiDDoS appliance's internal clock. Click *Change* to change the time or configure the FortiDDoS appliance to get the time from an NTP server. See "Setting the system time & date" on page 95.
System Uptime	Displays the time in days, hours, and minutes since the FortiDDoS appliance last started.
Effective HA Mode	Displays whether this FortiDDoS appliance is operating as a standalone appliance or as part of a high availability (HA) pair. A HA configuration synchronizes the configuration of a standby appliance with the active appliance. This value is determined by the current HA status of the appliance, not the current HA configuration. For HA pair configuration information, see "Topology for synchronizing the configuration of two FortiDDoS appliances" on page 64.
User Name	The name of the current logged in user.
Reboot	Reboots the appliance.

Setting name	Description
ShutDown	Shuts down the appliance.
Reset	Returns all settings to the factory defaults. See "Resetting profile data or the appliance configuration" on page 266.

See also

- Changing the FortiDDoS appliance's host name

License Information widget

The *License Information* widget on the dashboard displays FortiGuard IP Reputation Service registration and licensing information. See "FortiGuard IP Reputation Service" on page 130.

Click *Register* or *Renew* to go to the customer service and support web site.

Setting name	Description
Registration	Displays the account that registered the appliance with Fortinet Technical Support.
Hardware	Displays the hardware version of the registered appliance.
Firmware	Displays the firmware version of the registered appliance.
Enhanced Support	Displays whether the appliance is registered for FortiCare 8x5 Enhanced Support, which provides support during local business hours.
Comprehensive Support	Displays whether the appliance is registered for FortiCare 24x7 Comprehensive Support, which provides round-the-clock access to mission critical support services.
IP Reputation Service Contract Date	Specifies when the contract for FortiGuard IP Reputation Service expires.
IP Reputation Service Definitions	Specifies the build number of the FortiGuard IP Reputation Service definitions. To schedule updates for the definitions, see "FortiGuard IP Reputation Service" on page 130. Click *Update* to upload a .pkg file that contains IP addresses to add to the FortiGuard IP Reputation Service definition.

See also

- FortiGuard IP Reputation Service

CLI Console widget

The *CLI Console* widget enables you to enter CLI commands through the web UI, without making a separate Telnet, SSH, or local console connection.

The *CLI Console* widget requires that your web browser support JavaScript.

To use the console, click in the console area. FortiDDoS logs you in using the same administrator account you used to access the web UI. You can then type commands into the *CLI Console* widget. Alternatively, you can paste commands from or into the console.

The default prompt includes the appliance's host name, which is the appliance serial number (for example, `FI800B3913800022 #`). To change the host name, see "Changing the FortiDDoS appliance's host name" on page 184.

Depending on the account that you use to log in to the FortiDDoS appliance, you may not have access to all CLI commands. If your account does not have *Read-Write* access for a command, the message "Permission denied" is displayed. For more information, see "Permissions" on page 44.

Figure 28:*CLI Console* widget

See also

- Permissions
- Global web UI & CLI settings

SPP Attacks widget

A graph that displays the number of incoming and outgoing packets dropped or blocked for each Service Protection Profile (SPP) during the specified period. For more details, hover over a point on the graph.

To hide or display the count for a profile, click the profile name.

Use the Attack Graphs dashboard (*Log & Report* > *Attack Graphs* > *Attack Graphs)* to view this widget with additional options, such as isolating a particular profile or saving the information as a PDF file.

See also

- Attack Graphs dashboard

Event Log Console widget

The *Event Log Console* widget on the dashboard displays log-based messages related to system events. FortiDDoS displays messages related to attack activity in the DDoS Attack Log (*Log & Report* > *Log Access* > *DDoS Attack Log).*

Event logs help you track system events on your FortiDDoS appliance, such as firmware changes. For more detailed event information, click *Log & Report* > *Log Access* > *Event Log*. (See “Viewing log messages” on page 239.)

Event log messages can also be delivered by email, Syslog, or SNMP. For more information, see “Logging” on page 221, “Alert email” on page 241, and “SNMP traps & queries” on page 243.

Figure 29:*Event Log Console* widget

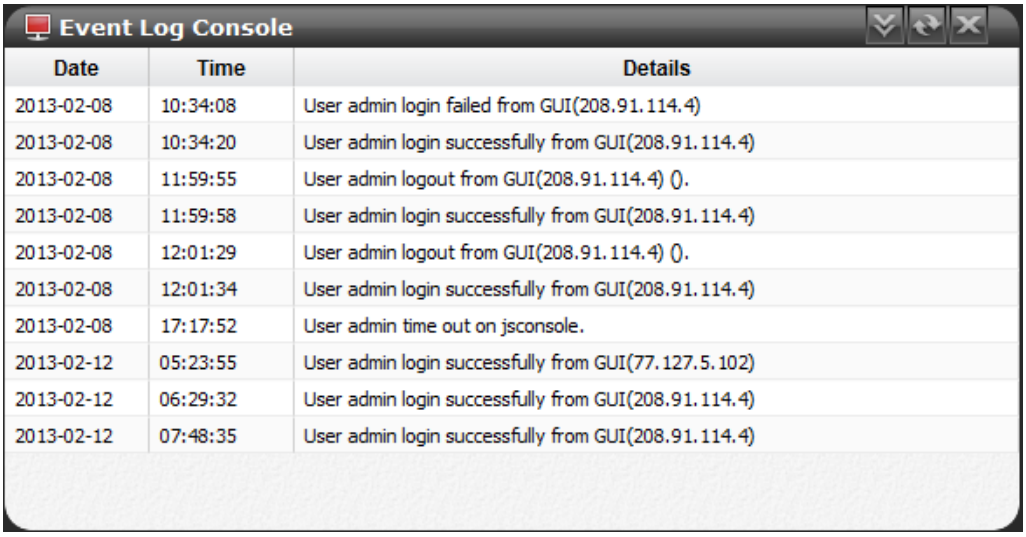

See also

- System event logs & logging

System Status widget

The *System Status* widget displays the status of the FortiDDoS appliance's link pairs and Service Protection Profiles (SPPs).

Item	Description
Port	Green icon - The port is physically connected the network. Red icon - The port has no physical connection to the network. Odd-numbered ports are LAN connections that have a corresponding even-numbered port, which is the associated WAN connection. (For example, Port 1 connected to the Ethernet and Port 2 connected to the Internet). Hover over the status icons to see additional information: port number, link status, speed, auto-negotiation, and medium (copper or fiber).
Bypass State	Used for copper-based (RJ-45) Ethernet connections only. Green icon - Normal operation (no bypass). Red icon - Link is operating in bypass mode. For more information on bypass features, see “External bypass switches for maintenance & failover” on page 55.
SPP ID	Green icon - Profile is operating in detection mode (monitoring traffic to generate statistics). Red icon - Profile is operating in prevention mode (dropping or blocking attacks as well as generating statistics). See “Setting FortiDDoS to detection mode” on page 116. The caution symbol indicates that the profile is not configured. See “Identifying IP addresses and subnets to protect (SPP creation)” on page 111. Hover over the status icons to see additional information: port number, ID, name, mode.

See also

- Configuring the network interfaces
- External bypass switches for maintenance & failover
- Setting FortiDDoS to detection mode
- Identifying IP addresses and subnets to protect (SPP creation)

Count of Unique Sources widget

Displays the number of unique sources flowing through FortiDDoS for the specified Service Protection Profile (SPP) and time period.

A spike in this graph indicates a possible DDoS attack.

To display the same information for an individual Service Protection Profile (SPP), use the Layer 3 traffic graph (*Monitor > Layer 3 > Count of Unique Sources)*.

See also

- Layer 3 graphs

System Resources widget

The *System Resources* widget on the dashboard displays CPU and memory usage.

The widget displays CPU and memory usage as a dial gauge and as a percentage of the usage for core processes only. CPU and memory usage for management processes (for example, for HTTPS connections to the web UI) is excluded.

Figure 30:*System Resources* widget

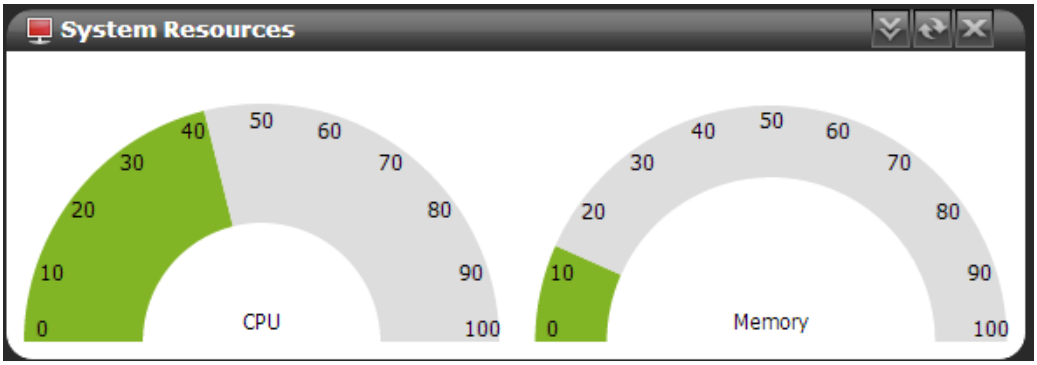

Disk usage is ***not*** displayed. To view disk space information, connect to the CLI and enter the following command:

```
diagnose hardware get sysinfo df
```

Traffic graphs

FortiDDoS starts learning traffic patterns the moment it is introduced in the network and it never stops learning. It continuously records traffic statistics that you can view in graphical form. FortiDDoS manages graph information in a round-robin fashion, meaning when you navigate to a graph or refresh it, FortiDDoS displays the most current information that is available for the specified time period. The graph information is always the previous 8-hours, day, week, month, or year.

You use the *Monitor* menu to display the traffic graphs. All the traffic is recorded every 5 minutes and the graphs display the highest per second rate that FortiDDoS observed during those 5 minutes. These rates give you an idea of traffic characteristics in your network.

Setting name	Description
Threshold	Displays the current configured minimum threshold for traffic parameters displayed in the graph.
Data Resolution	Displays the data resolution of the current graph. For example, for a period of 1 hour, 8 hours, or 1 day, the resolution is 5 minutes. For 1 week, it is 1 hour and for a month, 3 hours.

Setting name	Description
Refresh	Regenerates the graph using the latest statistics and the currently selected options.
Service Protection Profile	Specifies the protection profile to display.
Period	Specifies the time period to display.
Direction	Specifies the direction to display.

Dropped and blocked traffic statistics

Use the *Monitor* menu to display statistics for dropped and blocked traffic.

Dropped packets are packets that have been dropped due to header, rate, or state anomalies.

Blocked packets are packets that have been blocked due to access control lists (ACLs).

Like other traffic statistics, the dropped traffic is recorded every 5 minutes. FortiDDoS displays the actual counts of packets dropped and blocked during those 5 minutes.

These graphs give you an idea of anomalous traffic characteristics in your network.

Aggregate drops

The Aggregate Flood Drops and Aggregate ACL Drops graphs account for any packet that FortiDDoS drops or blocks.

FortiDDoS provides an aggregate drop graph for all layers and individual graphs for layers 3, 4, and 7.

Typical packet traffic graph

Figure 31 is a graph of packet traffic associated with a specific protocol (*Monitor* > *Specific Graphs* > *Protocols*).

Figure 31: Graph of inbound TCP traffic

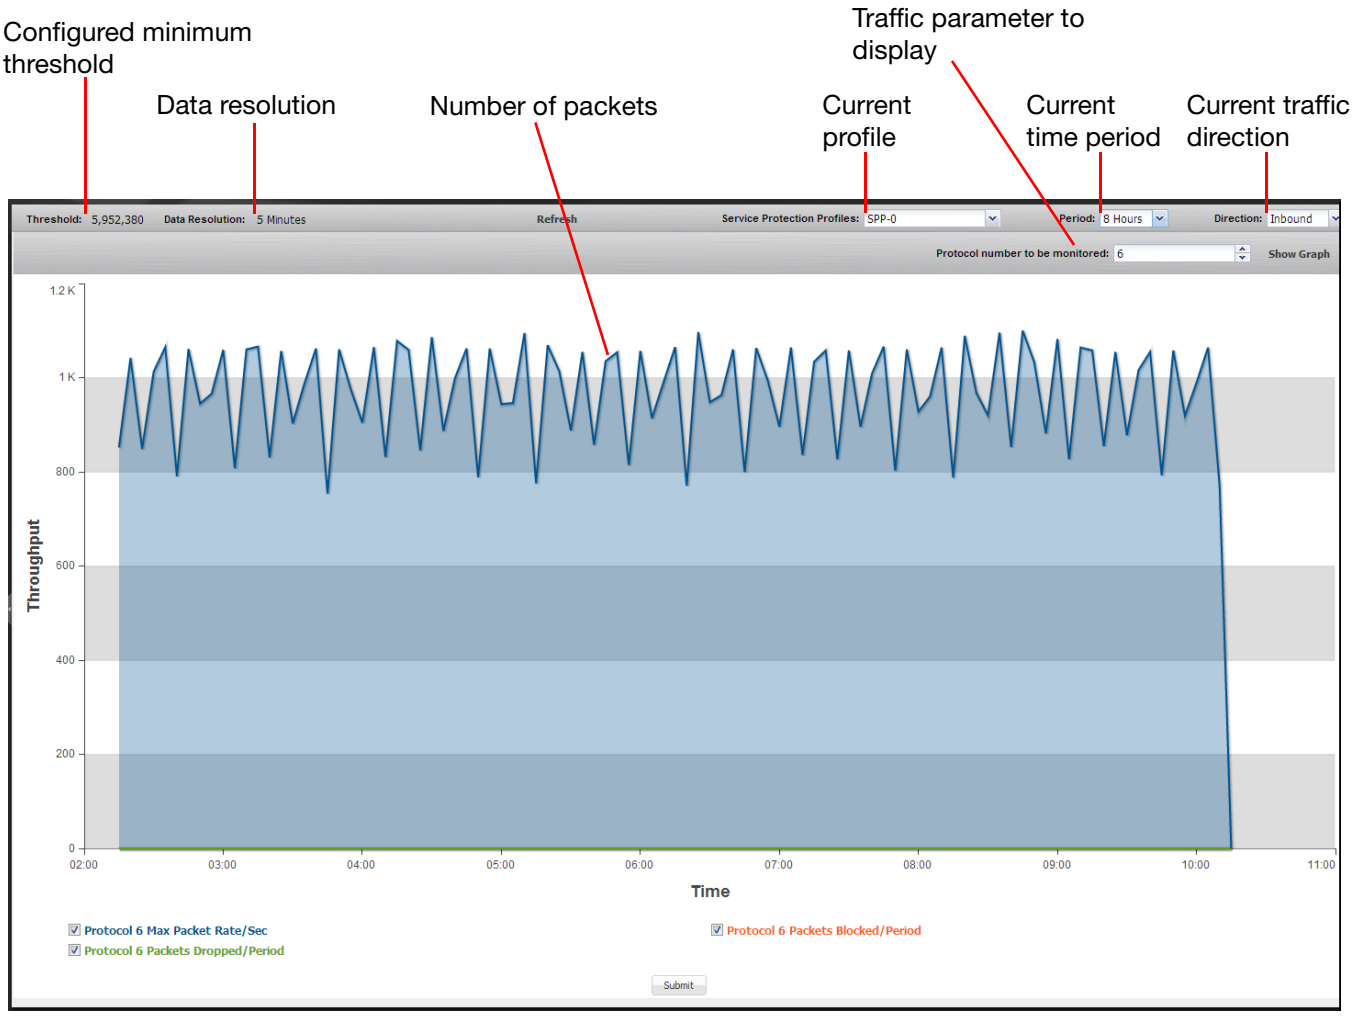

It shows the following information:

- The flow of TCP protocol packets through the FortiDDoS over an 8-hour period. The period is a round-robin representation, that is, it is always the previous hour, previous 8 hours, and so on.
- The inbound traffic rate. To view outbound traffic, for *Direction*, select *Outbound*.
- The horizontal axis shows the time scale while the vertical scale shows the throughput in packets per second.
- You can see that at this time of the day, traffic is steady and does not increase or decrease as the day progresses.
- Hover on the graph line to display detailed information for that point on the graph.

Figure 32: Tooltip information for point on graph line

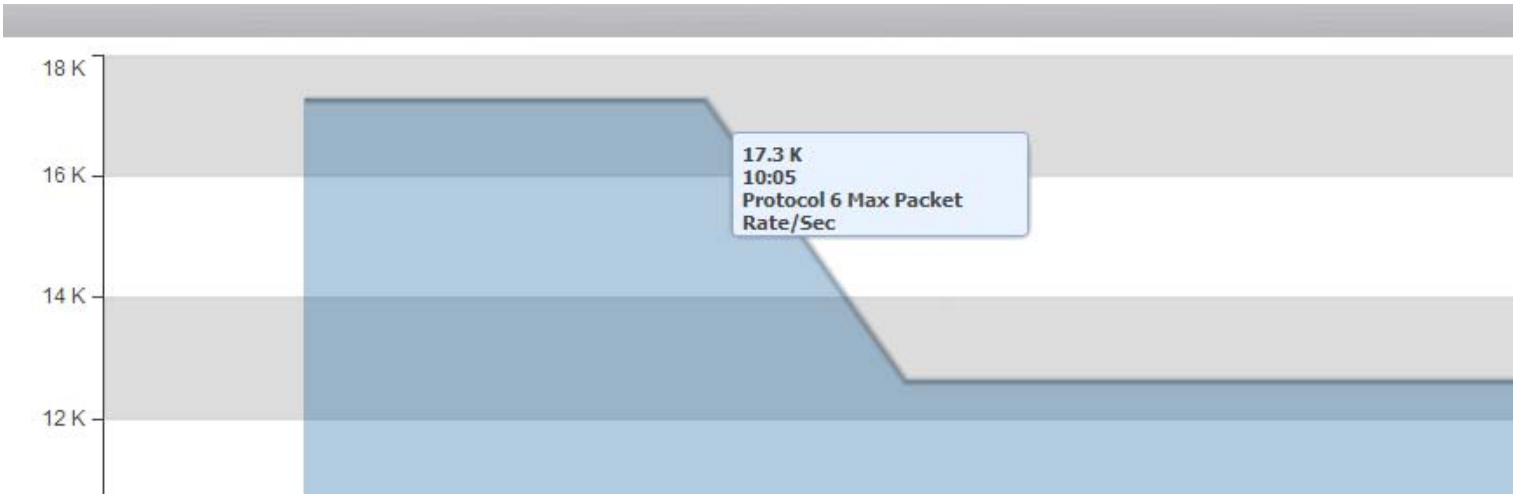

To change which graph lines are displayed, select or clear legend items at the bottom of the graph, and then click *Submit*.

Like most graphs, the resolution is 5 minutes. The reported traffic volumes are the highest packet rate within 1 second for the TCP protocol during the 5-minute sampling period.

Graph is correct only up to the last checkpoint period. For example, the 1-hour graph with a 5-minute resolution is correct only up to the time before the last 5 minutes. The data in the last 5 minutes may not yet have been registered. Therefore, a traffic peak in the last 5 minutes may not appear in the graphs immediately.

Many graphs include a line for the estimated threshold. For more information on how FortiDDoS generates and uses the estimated threshold, see “Continuous learning & adaptive threshold estimation” on page 23.

See also

- Working with graphs: Aggregate Flood Drops
- Working with graphs: Aggregate ACL Drops
- Traffic graphs for other counts

Working with graphs: Aggregate Flood Drops

Figure 33 illustrates the graph that summarizes packets dropped due to flood activity. FortiDDoS dropped these packets because they exceeded a threshold for a traffic parameter (rate anomaly) or contained header or state anomalies.

Figure 33: Aggregate Flood Drops graph

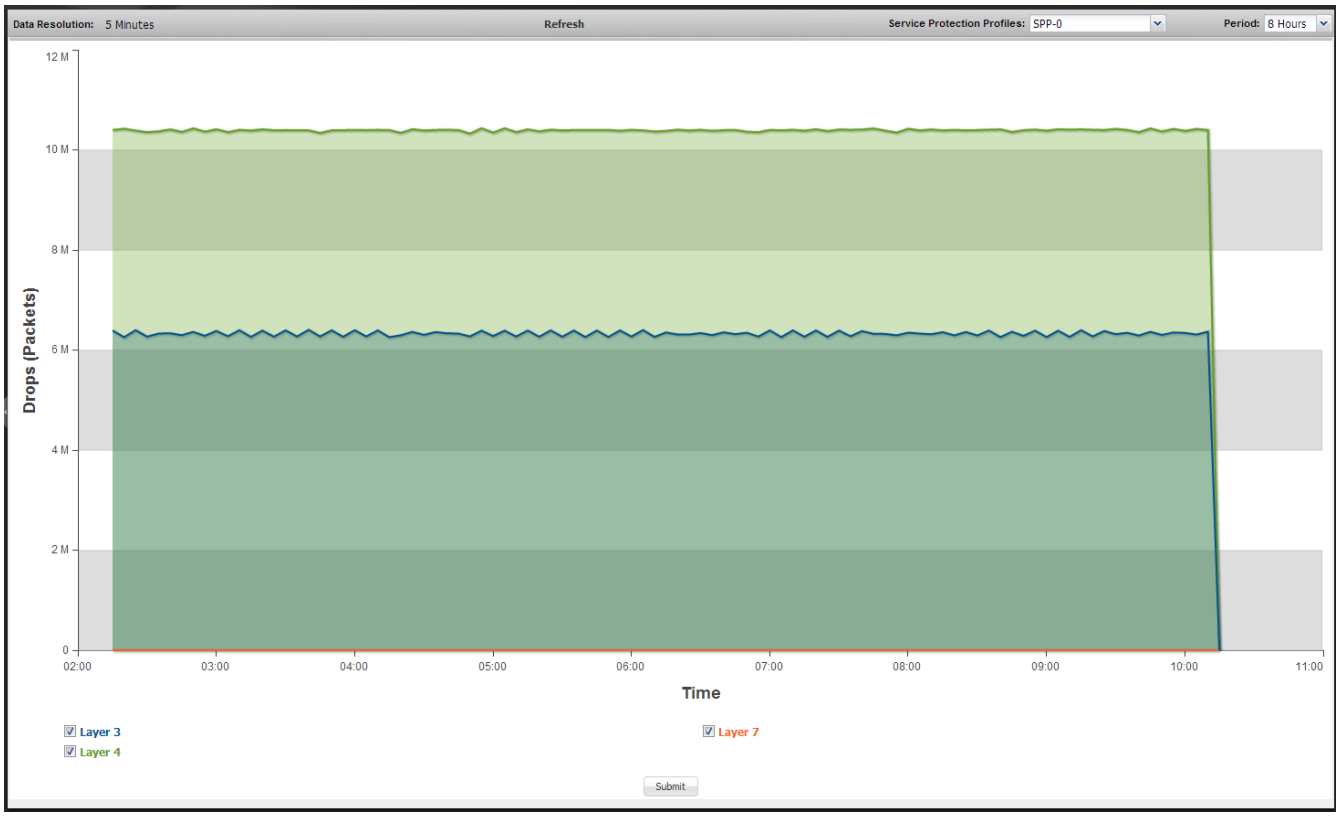

The graph illustrates the number of packets dropped due to layer 3, 4, and 7 parameters. It helps you to determine where to look for more information.

For example, if there is a significant volume of dropped packets for layer 3, go to *Aggregate Flood Drops* > *Layer 3,* which breaks down the dropped traffic by layer 3 parameters such as protocols and fragmented packets.

To narrow the investigation further, use the appropriate dashboards and graphs. For example, if FortiDDoS drops a high number of packets because they exceeded a protocol threshold, use the following dashboards to determine which protocols are associated with the drops:

- *Log & Report* > *Report Browse* > *Executive Summary*
- *Log & Report* > *Attack Graphs* > *Attack Graphs*

After you determine which specific protocols to check, go to *Monitor* > *Specific Graphs* > *Protocols* to retrieve graph information for individual protocols.

See also

- Traffic graphs
- Aggregate Flood Drops graphs

Working with graphs: Aggregate ACL Drops

Like the *Aggregate Flood Drops* graph, the information in the *Aggregate ACL Drops* graph summarizes the blocked packet traffic by layer and helps you to determine where to look for more information.

For example, if there is a significant volume of blocked layer 4 packets, go to *Aggregate ACL Drops* > *Layer 4*, which illustrates packets that FortiDDoS blocked because of layer 4 traffic parameters that are denied in the ACL for the specified protection profile.

Figure 34: Aggregate ACL Drops Layer 4 graph

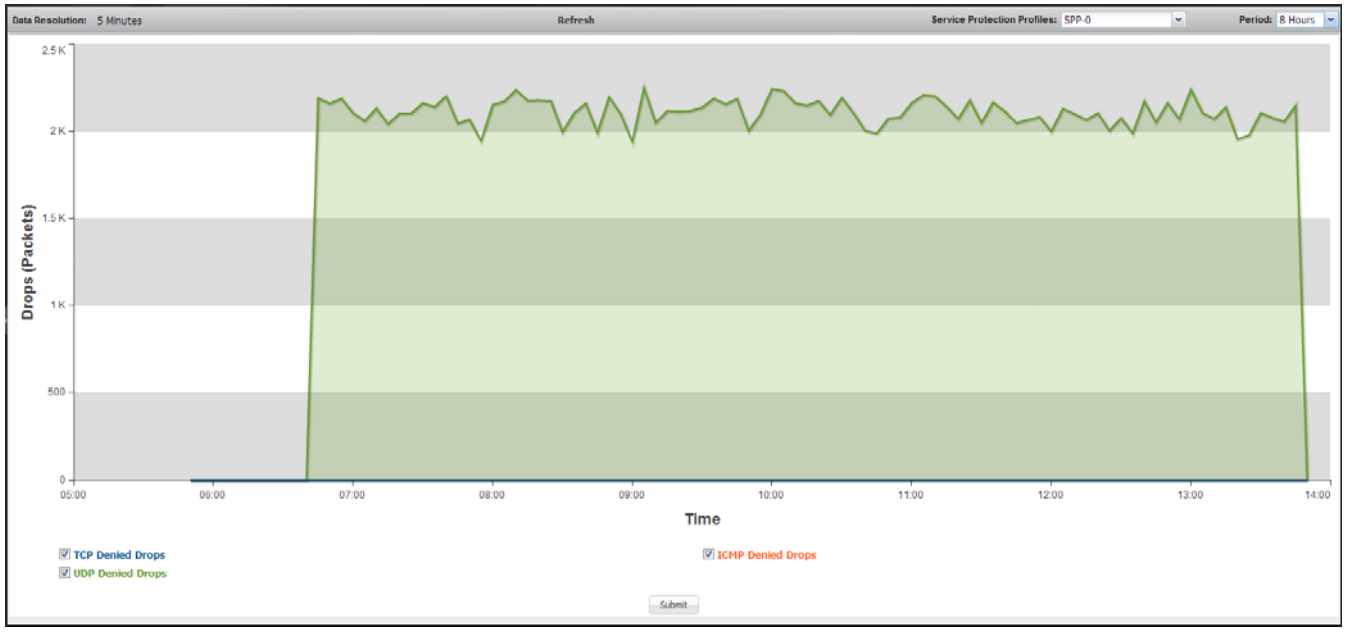

To narrow the investigation, use the appropriate dashboards and graphs. For example, if FortiDDoS blocks a high number of packets because they exceeded a UDP port threshold, use the following dashboards to determine which UDP ports are associated with the blocked packets:

- *Log & Report* > *Report Browse* > *Executive Summary*
- *Log & Report* > *Attack Graphs* > *Attack Graphs*

After you determine which specific ports to check, go to *Monitor* > *Specific Graphs* > *UDP Ports* to retrieve graph information for individual ports.

See also

- Traffic graphs
- Aggregate ACL Drops graphs

Traffic graphs for other counts

In addition to packet traffic, FortiDDoS graphs show statistics for other parameters, including the following information:

- Count of unique sources, Figure 35 (*Monitor* > *Layer 3* > *Count of Unique Sources*)
- Number of connections and number of entries in the TCP state table (shown in *Monitor* > *Layer 4* > *Established Connections*)
- Entries in the legitimate IP address table (shown in *Monitor* > *Layer 4* > *New Connections*)

Figure 35: Count of unique sources

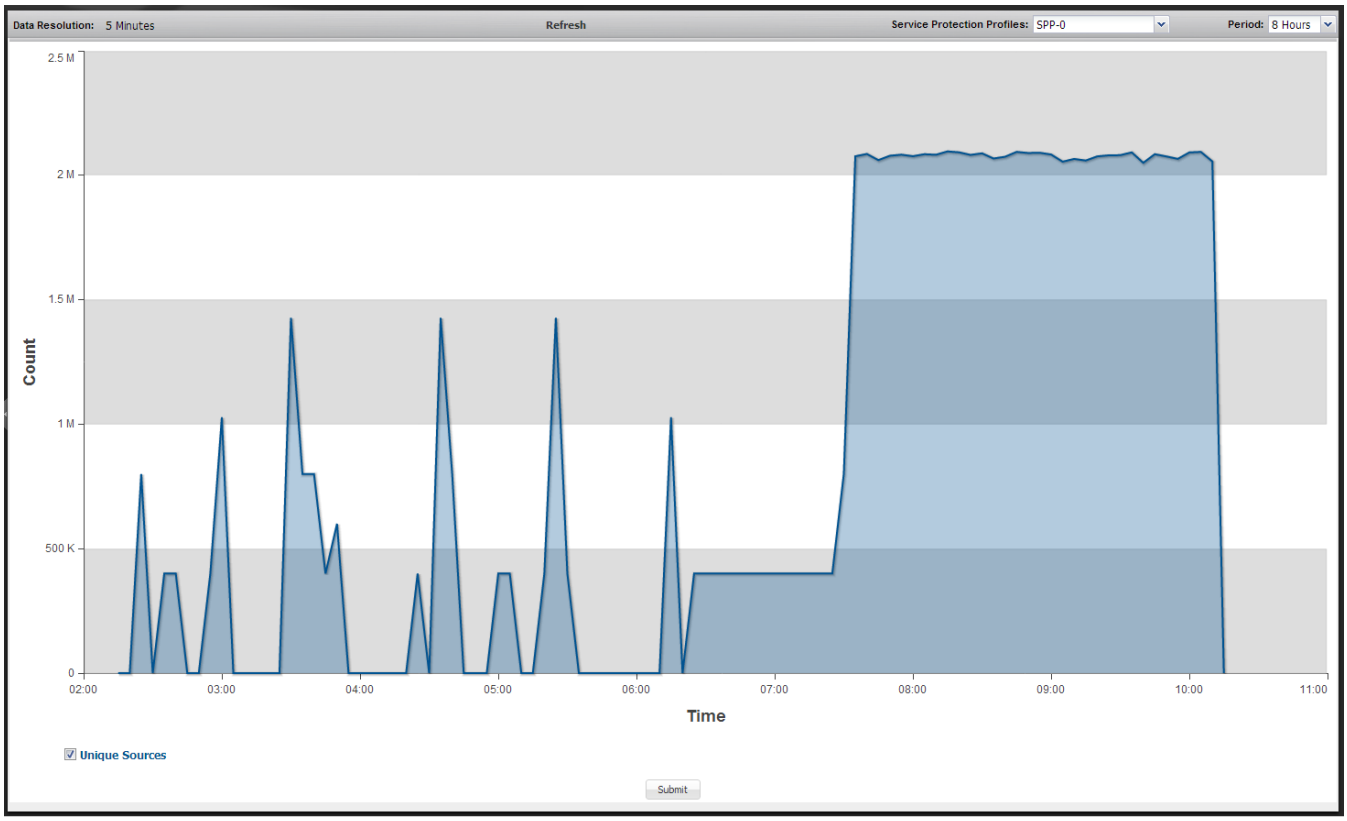

Since different time periods are handled with a different resolution, each graph is correct only up to the last checkpoint period. For example, the 1-hour graph with a 5-minute resolution is correct only up to the time before the last 5 minutes. The data in the last 5 minutes may not yet have been registered. Therefore, a traffic peak in the last 5 minutes may not appear in the graphs immediately. Similarly, for a 1-month graph with a data resolution of 3 hours, the data for the last 3 hours may not appear immediately.

See also

- Traffic graphs

Port Statistics graphs

The Port Statistics graphs allow you to view traffic passing through FortiDDoS in Mbps. These graphs illustrate the count of packets and bits travelling through a port pair.

The FortiDDoS appliance has 16 physical LAN and WAN ports, which are configured as port pairs. Odd-numbered ports are LAN connections that have a corresponding even-numbered port, which is the associated WAN connection. That is, Port 1/Port 2 behaves as LAN 1/WAN 1, Port 3/Port 4 as LAN 2/WAN 2, Port 5/Port 6 as LAN 3/WAN 3, and so on.

Unlike most of the FortiDDoS graphs, which have a resolution of 5 minutes, the Port Statistics graphs have a resolution of 30 seconds.

Viewing the inbound and outbound traffic that the physical LAN and WAN ports receive can increase your understanding of overall traffic patterns and allows you to view FortiDDoS attack mitigation at work.

The graph in Figure 36 illustrates inbound traffic for a port pair in Kbps. FortiDDoS is in the process of intercepting attack traffic, so the count of packets for the ingress port (port 2 or WAN port) is visibly higher than the count at the egress port (port 1 or LAN port).

Figure 36: Example port statistics graph showing packets in Kbps

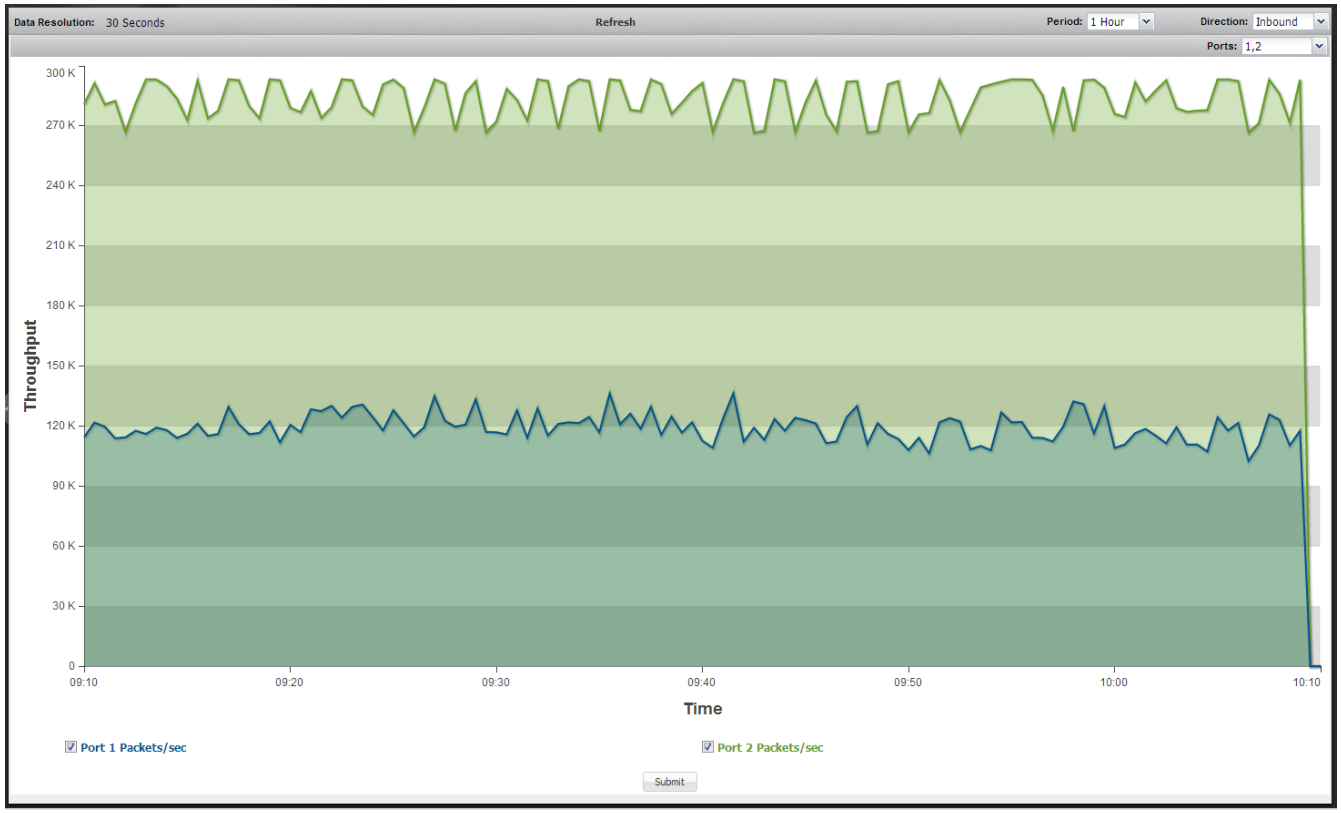

To view the packet or byte count for a link pair

1. Do one of the following:
   - Click *Monitor* > *Port Statistics* > *Packets*.
   - Click *Monitor* > *Port Statistics* > *Bits*.
2. Select the time period and traffic direction for the graph.
3. For *Ports*, select the appropriate link pair.

See also

- Traffic graphs

My Graphs

The *My Graphs* graphs can display statistics for all the traffic parameters where the number of items you can configure or monitor is more than a handful. For example, the *My Graphs* > *Protocols* graph can display any of the 256 independent protocols you set thresholds for and monitor. In most cases, there are only a few protocols (such as TCP, UDP, ICMP) that you are interested in, since most Internet traffic is associated with these protocols.

FortiDDoS provides the following three options for adding items to *My Graphs*, all of which you access using *Protection Profiles* > *My List*. You configure the *My Graphs* items for each Service Protection Profile (SPP) individually:

- **My List Details** — Allows you to manually add an item to *My Graphs* by specifying its type and value.
- **Autoconfigure** — Adds TCP, UDP, and ICMP ports, protocols, URLs, and HTTP header field items to *My Graphs* based on a traffic report you generate. It adds the top 5 entries in the report for each graph. For example, it adds the 5 most active TCP ports to *My Graphs* > *TCP Ports*.
- **Set to Defaults** — Adds items to *My Graphs* that are useful in most environments.
  - Protocols: TCP (6), UDP (17)
  - TCP Ports: 22, 23, 25, 80, 443
  - UDP Ports: 53
  - ICMP Types/Codes: 0, 2048
  - URLs: 15233 (/index.html), 2502 (/index.htm), 453 (index.php)

To manually add an item to My Graphs via the web UI

1. Click *Protection Profiles* > *My List* > *My List Details*.
2. For *Service Protection Profile*, select the profile to add the *My List* item to.
3. For *Type*, select the type of item you want to add.
4. Click *Add*.
5. Enter a name and value for the item.
6. Click *Save*.

To add items to My Graphs using Autoconfigure via the web UI

1. Generate a traffic statistics report for the appropriate service protection profile (SPP) and time period.

   For detailed instructions, see “Generating and reviewing a traffic statistics report” on page 144.
2. Click *Protection Profiles* > *My List* > *Autoconfigure*.
3. Selelct the appropriate SPP and time period, select *Enable*, and then click *Save*. Click *OK* to confirm the change.
4. Click *Protection Profiles* > *My List* > *My List Details* to review the items that FortiDDoS added.

To add popular items to My Graphs via the web UI (Set to Defaults)

1. Generate a traffic statistics report for the appropriate service protection profile (SPP) and time period.

   For detailed instructions, see “Generating and reviewing a traffic statistics report” on page 144.
2. Click *Protection Profiles* > *My List* > *Autoconfigure*.
3. Selelct the appropriate SPP and time period, select *Enable*, and then click *Save*. Click *OK* to confirm the change.
4. Click *Protection Profiles* > *My List* > *My List Details* to review the items that FortiDDoS added.

To add a protocol to My Graphs via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp mylist-protocol
    edit <threshold_name>
    set protocol <protocol_int>
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `<threshold_name>` specifies the name of the protocol in the list of items
- `<protocol_int>` specifies the protocol to add to the *My Graph* traffic graphs

To add a TCP or UDP port to My Graphs via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp {mylist-tcp-port | mylist-udp-port}
    edit <threshold_name>
    set protocol <port_int>
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `{mylist-tcp-port | mylist-udp-port}` specifies whether the item is a TCP or UDP port
- `<threshold_name>` specifies the name of the port in the list of items
- `<port_int>` specifies the port to add to the *My Graph* traffic graphs

To add an ICMP type and code to My Graphs via the CLI

Enter the following commands:

```
edit <spp_name>
  config ddos spp mylist-icmp-type-code
    edit <threshold_name>
    set icmp-type-code <type_code_int>
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `<threshold_name>` specifies the name of the ICMP type and code in the list of items
- `<type_code_int>` specifies an ICMP type and code using the following formula: 256 * ICMP type number + ICMP code

To add an HTTP URL or header field to My Graphs via the CLI

Enter the following commands:

```
config spp
  edit <spp_name>
    config ddos spp {mylist-http-url | mylist-http-host |
    mylist-http-referer | mylist-http-cookie |
    mylist-http-user-agent}
      edit <my_graphs_name>
        set {url | host | referer | cookie | user-agent} test
end
```

where:

- `<spp_name>` is the name of the Service Protection Profile (SPP)
- `{mylist-http-url | mylist-http-host | mylist-http-referer | mylist-http-cookie | mylist-http-user-agent}` specifies whether to add a URL or HTTP header field to the *My Graphs* traffic graphs
- `<my_graphs_name>` specifies the name of the URL or header field in the list of items
- `{url | host | referer | cookie | user-agent}` specifies the type of item added to the *My Graphs* traffic graphs

See also

- Traffic graphs

Specific Graphs

These graphs display the packet rate in both inbound and outbound directions by port, URL, header field or other characteristic.

For example, *Specific Graphs* > *Protocols* displays the rate of packets per second and dropped and blocked packet count for the protocol you specify.

Use *My Graphs* to access customized lists of specific graphs.

For the following graphs, to specify the information to display, you enter a specific value: URLs, Host, Referer, Cookie, User Agent.

Because there are an infinite number of possible URLs and Host, Referer, Cookie, and User-Agents header field values, FortiDDoS assigns these items to a hash index. The graph legend displays the hash index for the value that you specify.

FortiDDoS provides the following specific graphs:

- Protocols
- TCP Ports
- UDP Ports
- ICMP Types/Codes
- URLs
- Hosts
- Referers
- Cookies
- User Agents

See also

- Traffic graphs

Aggregate Flood Drops graphs

FortiDDoS provides the following flood drops graphs:

- Aggregate Flood Drops graph (all layers)
- Layer 3 Aggregate Flood Drops graph
- Layer 4 Aggregate Flood Drops graph
- Layer 7 Aggregate Flood Drops graph

Aggregate Flood Drops graph (all layers)

Displays a summary of all dropped packets from flood events by layer.

For an example of this type of graph, see “Working with graphs: Aggregate Flood Drops” on page 205.

See also

- Traffic graphs
- Working with graphs: Aggregate Flood Drops

Layer 3 Aggregate Flood Drops graph

This graph displays the following dropped packet information:

- Protocols

  Dropped and blocked packet count in both inbound and outbound directions for all protocols. To view statistics for a specific protocol, use *Specific Graphs* > *Protocols* or add an item to *Protection Profiles* > *My List* > *My List Details*.
- Fragmented Packets

  Packets that are blocked or dropped due to fragmentation error in both inbound and outbound directions. For a more detailed fragmented packets graph, click *Monitor* > *Layer 3* > *Fragmented Packets*.
- L3 Anomalies

  Packets with one or more of layer 3 header anomalies, including the following anomalies:
  - IP version 4 header checksum errors
  - Same IP address for source and destination (indicates a LAND attack)
  - Loopback address spoofing errors
  - Local Host — Illegal use of IPv4/IPv6 loopback address for source, destination, or both

  For a detailed layer 3 anomaly drop graph, click *Monitor* > *Anomaly Drops* > *Layer 3 Anomaly Drops*.
- Source Flood

  Dropped and blocked packets from over-active sources. Dropped sources are sources that FortiDDoS has dropped due to floods while the blocked sources correspond to ACL entries in the layer 3 sources.

Use the following dashboards and graphs for more information about over-active sources:

- *Log & Report* > *Report Browse* > *Executive Summary*
- *Log & Report* > *Attack Graphs* > *Attack Graphs*
- *Monitor* > *Layer 3* > *Most Active Source*

For more detailed ACL drops graphs, click *Aggregate ACL Drops* > *Layer 3*, *Layer 3* > *Dark Address Scan*, and *Layer 3* > *Denied Countries*.

- Misc Source Flood

  Drops due to hash attack and memory limitations. Because the source table is a large, dynamic table with up to 2 million entries, it can be attacked using hash attacks or may run out of memory for pointers.

  More detailed hash attack and memory graphs are located under *Monitor* > *Hash Attack Drops* and *Monitor* > *Out of Memory Drops*.

- Destination Flood

  Count of dropped and blocked packets from over-active or blocked destinations. Dropped destinations are sources that FortiDDoS has dropped due to floods while the blocked destinations correspond to ACL entries.

  For a more detailed graph of over-active destinations, click *Monitor* > *Layer 3* > *Most Active Destination*.

  For more detailed ACL drops graphs, click *Aggregate ACL Drops* > *Layer 3*, *Layer 3* > *Dark Address Scan*, and *Layer 3* > *Denied Countries*.

- Misc Destination Flood

  Count of drops due to hash attack and memory limitations. Because the destination table is a large, dynamic table with up to 2 million entries, it can be attacked using hash attacks or may run out of memory for pointers.

  More detailed hash attack and memory graphs are located under *Monitor* > *Hash Attack Drops* and *Monitor* > *Out of Memory Drops*.

- Dark Address Scan

  Count of packets dropped because the associated network is denied in the global ACL.

- URL Source Tracking

  Count of packets dropped because traffic has exceeded the threshold for a specific URL.

See also

- Traffic graphs
- Working with graphs: Aggregate Flood Drops

Layer 4 Aggregate Flood Drops graph

This graph displays lines for the following drop information:

- SYN Packets

  Dropped and blocked SYN packets in both inbound and outbound directions. SYN packets are sent to open a new TCP connection.

  FortiDDoS provides more detailed dropped and blocked SYN packet graphs under *Monitor* > *Layer 4*.

- L4 Anomalies

Packets with one of the following layer 4 header anomalies:

- TCP, UDP and ICMP header checksum errors
- Invalid TCP flag combinations
- Other layer 4 header anomalies such as incomplete packet

For more detailed layer 4 anomaly graphs, click *Anomaly Drops* > *Layer 4 Header Anomalies*.

- TCP Ports and UDP Ports

  Dropped and blocked packets in both inbound and outbound directions for all TCP and UDP port. To view statistics for a specific port, under *Specific Graphs*, click *TCP Ports* or *UDP Ports*.

- ICMP Types/Codes

  Dropped and blocked packet count in both inbound and outbound directions for every ICMP type and code. To view statistics for a specific type and code, use *Specific Graphs* > *ICMP Types/Codes*.

- Misc. Connection Flood

  Because the TCP connection table is a large dynamic table with up to 1 million entries, it can be attacked using hash attacks or may run out of memory for pointers. This line illustrates the drops due to hash attack and memory limitations.

  More detailed hash attack and memory graphs are located under *Monitor* > *Hash Attack Drops* and *Monitor* > *Out of Memory Drops*.

  In addition, this line includes the count of packets with anomalous TCP packets which do not meet TCP state transition rules. These anomalies are:

  - TCP window size violations
  - Foreign TCP packets
  - TCP state transition anomalies
  - Forward window size violation - (invalid sequence number).
  - Reverse window size violation - (invalid sequence number).
  - Drop due to non-SYN - A packet for a connection has been received that has not been established with a SYN exchange.

  More detailed TCP state anomalies graphs are located under *Anomaly Drops* > *TCP State Anomalies*.

- Zombie Flood

  Count of packets dropped because traffic has exceeded the New Connections threshold, which sets a limit for legitimate IPs. FortiDDoS assumes a zombie flood is underway when the number of allowed legitimate IP addresses during a SYN flood exceeds a set threshold. These packets indicate that non-spoofed IP addresses are creating a distributed DoS attack by generating a large number SYN packets.

  For a more detailed graph, go to *Monitor* > *Layer 4* > *New Connections*.

- SYN Packets Per Source Flood

  Dropped and blocked packets in both inbound and outbound directions that were dropped due to a single source sending too many SYN packets.

  For a more detailed graph, click *Layer 4* > *SYN Per Source*.

- Excessive Concurrent Connections Per Source

  Dropped and blocked packets in both inbound and outbound directions that were dropped because a single source built too many TCP concurrent connections.

  For a more detailed graph, click *Layer 4* > *Connection Per Source*.

- Excessive Concurrent Connections Per Destination

Dropped and blocked packets in both inbound and outbound directions that were dropped because a single destination was overloaded with too many idle connections.

For a more detailed graph, click *Layer 4* > *Connection Per Destination*.

- TCP Packets Per Destination

  Count of packets in both inbound and outbound directions that were dropped because traffic exceeded one of the following thresholds: SYN, FIN, ACK, or RST packets per destination, or established connections per destination. For individual graphs for these counts, see the *Layer 4* graphs.

See also

- Traffic graphs
- Working with graphs: Aggregate Flood Drops

Layer 7 Aggregate Flood Drops graph

This graph displays lines for the following drop information, in both inbound and outbound directions:

- Method Flood

  Dropped packet count for all HTTP Method values. To view statistics for a specific method, use *Layer 7* > *Protocols*.

- HTTP Anomalies

  Dropped and blocked packet count for anomalous layer 7 headers.

- URL, Host, Use Agent, Cookie, and Referer Flood

  Dropped and blocked packet count for URLs and HTTP header field values. To view statistics for a specific URL or header field, under *Specific Graphs*, select the appropriate graph.

- Invite Per Source Drop

  Dropped and blocked packets in both inbound and outbound directions that were dropped due to a single source sending too many SIP INVITE requests.

  To view more detailed statistics, use *Layer 7* > *Invites*.

- Register Per Source Drop

  Packets in both inbound and outbound directions that were dropped due to a single source sending too many SIP REGISTER requests. To view more detailed statistics, use *Layer 7* > *Registers*.

- Concurrent Invites Per Source Drop

  Packets in both inbound and outbound directions that were dropped because a single source built too many concurrent INVITE sessions. To view more detailed statistics, use *Layer 7* > *Concurrent Invites*.

Alternative and detailed views of some of this information are available in the *Specific Graphs*, *Aggregate ACL Drops*, *Anomaly Drops*, and *Layer 7* graphs.

You can also generate graphs for specific parameters using *Protection Profiles* > *My List* > *My List Details*. For example, to monitor a specific set of URLs, add the URLs to *My List Details*. FortiDDoS adds statistics for each URL to *My Graphs* > *URLs*.

See also

- Traffic graphs
- Working with graphs: Aggregate Flood Drops

Aggregate ACL Drops graphs

For more information on denying traffic using access control lists (ACLs), see “Access control lists (ACLs)” on page 119.

Aggregate ACL Drops graph (all layers)

Displays counts of all packets dropped because of ACL configuration by layer.

Layer 3 Aggregate ACL Drops graph

Displays counts of all packets dropped because of ACL configuration for layer 3 traffic parameters.

Layer 4 Aggregate ACL Drops graph

Displays counts of all packets dropped because of ACL configuration for layer 4 traffic parameters.

Layer 7 Aggregate ACL Drops graph

Displays counts of all packets dropped because of ACL configuration for layer 7 traffic parameters.

See also

- Traffic graphs
- Working with graphs: Aggregate ACL Drops
- Access control lists (ACLs)

Anomaly Drops graphs

Layer 3 Anomaly Drops graph

This graph displays the count of packets that FortiDDoS dropped or blocked due to the following layer 3 header anomalies:

- IP header checksum errors
- Source and destination addresses are the same (LAND attack)
- Source or destination address is the same as the Local Host (loopback address spoofing)
- IP version other than 4 or 6
- Header length less than 5 words
- End of packet (EOP) before 20 bytes of IPV4 Data
- Total length less than 20 bytes
- EOP comes before the length specified by Total length
- End of Header before the data offset (while parsing options)
- Length field in LSRR/SSRR option is other than (3+(n*4)) where n takes value greater than or equal to 1
- Pointer in LSRR/SSRR is other than (n*4) where n takes value greater than or equal to1
- For IP Options length less than 3

Layer 4 Header Anomalies drop graph

This graph displays the count of packets that FortiDDoS dropped or blocked due to the following layer 4 header anomalies:

- TCP, UDP, and ICMP header checksum errors
- Invalid TCP flag combinations
- Other header anomalies, such as incomplete packet
- Urgent flag is set then the urgent pointer must be non-zero
- SYN or FIN or RST is set for fragmented packets
- Data offset is less than 5 for a TCP packet
- End of packet is detected before the 20 bytes of TCP header
- EOP before the data offset indicated data offset
- Length field in Window scale option other than 3 in a TCP packet
- Missing UDP payload
- Missing ICMP payload

TCP State Anomalies drop graph

The count of packets that FortiDDoS dropped or blocked due to the following TCP state anomalies:

- Forward or Reverse Transmission Not Within Window

  Packets are outside the receiver's TCP or UDP window.
- Session Traffic Threshold Cross
- TCP State Transition

Packets violate the TCP Protocol state transition rules or sequence numbers.

- Foreign packet

  Packets do not belong to a known TCP connection. For example, FortiDDoS detects a packet for a connection that has not been established with a SYN exchange. FortiDDoS can store up to 1 million TCP connections at a time in its internal memory.

HTTP Header Anomalies drop graph

The count of packets that FortiDDoS dropped or blocked due to the following header anomalies:

- Unknown method in HTTP header
- Invalid HTTP version

See also

- Traffic graphs

Hash Attack Drops and Out of Memory Drops graphs

The count of packets that FortiDDoS dropped because of a problem with internal FortiDDoS logic or memory. If these graphs report any dropped traffic, contact Fortinet for assistance.

See also

- Traffic graphs

Layer 3 graphs

- Most Active Source and Most Active Destination

  The count of packets sent to and received by the most active source or destination address for the specified period. Also the count of packets associated with the address that FortiDDoS dropped or blocked.

  Note that this is not necessarily a graph of a single address over time, since FortiDDoS samples the maximum packet value for all sources and destinations each second.

  The Most Active Source graph is useful for detecting single-source attacks, in which a source is sending an abnormally high number of packets.
- Count of Unique Sources

  The instantaneous count of unique sources flowing through FortiDDoS. A spike in this graph indicates a possible DDoS attack.
- Fragmented Packets

  The count of fragmented packets flowing in both inbound and outbound directions.
- Dark Address Scan

  This graph displays the count of packets that FortiDDoS has blocked because their IP address is blocked by the global ACL configuration.
- Denied Countries

  This graph displays the count of packets that FortiDDoS has blocked because their geographic location is blocked by the global ACL configuration.

See also

- Traffic graphs

Layer 4 graphs

- SYN Packets

  SYN packets flowing through a Service Protection Profile in both inbound and outbound directions.
- SYN Per Source

  The traffic rate in both inbound and outbound directions for the most active source sending TCP SYN packets among all the sources, recorded every 5 minutes. A spike in this graph shows a possible SYN attack from a single source or a few limited sources.
- SYN Per Destination

  The traffic rate in both inbound and outbound directions for the most active destination receiving TCP SYN packets among all the destinations in the profile, recorded every 5 minutes. A spike in this graph shows a possible SYN attack on a single destination or a few limited destinations.
- Connection Per Source

  The count of concurrent connections for the busiest source. A spike in this graph shows that a single source may be trying to establish too many connections.
- Connection Per Destination

  The count of concurrent connections for the busiest destinationconnection per. A spike in this graph shows a possible DDoS attack on a single destination.
- ACK Per Destination

  The traffic rate in both inbound and outbound directions for the most active destination receiving TCP ACK packets among all the destinations in the profile, recorded every 5 minutes. A spike in this graph shows a possible ACK attack on a single destination or a few limited destinations.
- RST Per Destination

  The traffic rate in both inbound and outbound directions for the most active destination receiving TCP RST packets among all the destinations in the profile, recorded every 5 minutes. A spike in this graph shows a possible RST attack on a single destination or a few limited destinations.
- FIN Per Destination

  The traffic rate in both inbound and outbound directions for the most active destination receiving TCP FIN packets among all the destinations in the profile, recorded every 5 minutes. A spike in this graph shows a possible FIN attack on a single destination or a few limited destinations.
- ESTAB Per Destination

  Traffic statistics in both inbound and outbound directions for the destination with the most established TCP connections, recorded every 5 minutes. A spike in this graph shows a possible connection establishment attack.
- New Connections

  The traffic rate in both inbound and outbound directions for new TCP connections established every second. A spike in this graph shows a possible concerted DoS or DDoS attack. You set a threshold for this type of traffic using the *established-connections* threshold.

This graph also illustrates the number of entries in the Legitimate IP address table.

- Established Connections

  The count of maximum instantaneous TCP connections that have completed three-way handshake. This is recorded every 5 minutes. A spike in this graph shows a possibility of a DoS or DDoS attack.

  Also illustrates the estimated threshold for established connections and the number of entries in the TCP state table.

  If the values for the number of entries in the TCP state table are significantly higher than those for established connections, it shows a possible SYN flood attack. FortiDDoS uses all entries in the connection table, including the half-open connections, to generate this graph.

See also

- Traffic graphs

Layer 7 graphs

- HTTP Methods

  The traffic rate for the HTTP method and direction you specify. The following methods are available:

  - GET
  - HEAD
  - OPTIONS
  - TRACE
  - POST
  - PUT
  - DELETE
  - CONNECT
- Invites

  The traffic rate in the direction you specify for the most active source sending SIP INVITE packets among all the sources, recorded every 5 minutes. A spike in this graph shows a possible SIP INVITE attack from a single source or a few limited sources.
- Registers

  The traffic rate in the direction you specify for the most active source sending SIP REGISTER packets among all the sources, recorded every 5 minutes. A spike in this graph shows a possibility of SIP REGISTER attack from a single source or a few limited sources.
- Concurrent Invites

  The count of concurrent SIP INVITE sessions for the busiest source, in the direction you specify. A spike in this graph can indicate that a single source is trying to establish too many sessions.

See also

- Traffic graphs

Logging

FortiDDoS provides two types of log messages: DDoS attack events and system events.

Both types of messages are the source of information for many types of reports. System event log messages are also the source of information for alert email.

Both types of logging are enabled by default, but you can disable system event logging if required.

See also

- DDoS Attack Log and DDoS Subnet Attack Log
- System event logs & logging
- Viewing log messages

DDoS Attack Log and DDoS Subnet Attack Log

- The `admin` administrator and administrators for whom the *System Admin* option is enabled can view reports for an individual SPP or all SPPs.
- The DDoS Attack Log displays all events for the specified SPP. FortiDDoS updates the DDoS Attack Log every few seconds.
- The DDoS Subnet Attack Log dislays events associated with a specific SPP policy. For example, the SPP policy that specifies the IP range for a specific group of web servers.
- FortiDDoS updates the DDoS Subnet Attack Log every five minutes and it displays a drop count for each event that is the total number of dropped packets in that period of time.
- By default, the DDoS Attack Log displays the events with the fields TimeStamp, SPP, DIR (direction), Event Type, and Drop Count. The DDoS Subnet Attack Log also displays a Reason Code value instead of Event Type, by default.
- Additional information for the selected event is displayed below the list of attack events, including Protocol, Source IP, Destination IP, Destination Port, ICMP Type/Code, Event Detail.
- Fields can display ”-” (hyphen) when there is no valid data for the fields.
- *Protocol*, *Destination Port*, and *ICMP Type/Code* use a number/name format. If the name cannot be determined, a ”-” (hyphen) is displayed.
- The DDoS Attack Log contains a maximum of 1 million events. If the number of events exceeds 1 million, FortiDDoS deletes the 200,000 oldest events.

The DDoS Attack Log event types are summarized in Table 11.

You can filter log events based on criteria such as category and time and date. For information on filtering the log, see “Filtering log messages” on page 240.

Table 11: DDoS Attack Log event categories

Layer	Category of event	Traffic parameter	Description
3	Rate Flood	Protocol	Excessive number of packets with specified protocol in the header.
		Fragment	Excessive number of fragmented IP packets. Some attacks use excessive fragmented packets to attempt to disable communications on the target.
		SYN per Source	Excessive number of TCP SYN packets from a single source.
		Tracked Source	
		TCP SYN, ACK, FIN, RST Per Destination	Excessive number of SYN, ACK, FIN, or RST packets to a destination.

Table 11: DDoS Attack Log event categories

Layer	Category of event	Traffic parameter	Description
3	ACL	Protocol	The global ACL or the ACL for a Service Protection Profile (SPP) can block the following: • Packets with a specified protocol • Fragmented IP packets • Packets with a specified source or destination IP address • Traffic from a specified country or region, or any anonymous proxy or satellite provider
		Fragment	
		Source	
		Destination	
		Geolocation	
	Scan Events	Dark Address Scan	The ACL can block packets associated with a specified network.
	Header Anomaly Events	IP header anomaly	An IP packet detected with one or more of the following anomalies: • Invalid IP header checksum • IP version other than IPv4 or IPv6 • Identical source and destination IP addresses (LAND attack) • Source or destination address is the local host (loopback address spoofing) • Invalid header length (less than 5 words) • EOP (End of Packet) before 20 bytes of IPv4 data • Total length less than 20 bytes • EOP comes before the length specified by Total Length • End of Header before the data offset (while parsing options) • Length field in the LSRR or SSRR IP option is other than (3+(n*4)) where n is a value greater than or equal to 1 • Pointer in the LSRR or SSRR IP option is other than (n*4) where n is a value greater than or equal to1

Table 11: DDoS Attack Log event categories

Layer	Category of event	Traffic parameter	Description
3	Internal Reason	Excessive hash collisions	An internal event occurs when there are excessive hash collisions, the system is out of memory, or there is some other problem with internal FortiDDoS logic or memory. If FortiDDoS drops packets for this reason, contact Fortinet for assistance.
		Out of memory	
		Source tracking: Hash attack	
		Source tracking:Out of memory	
		Destination Tracking: Hash attack	
		Destination tracking: Out of memory	

Table 11: DDoS Attack Log event categories

Layer	Category of event	Traffic parameter	Description
4	Rate Flood	SYN	Excessive TCP connection requests. Because TCP requires a three-way handshake to establish a connection, attackers that begin but do not finish the handshake process can absorb all resources reserved for legitimate users. For more information, see “SYN flood and zombie flood prevention” on page 135.
		SYN Flood From Source	Excessive number of SYN packets from a source address.
		SYN Flood to Destination	Excessive number of SYN packets to a destination address.
		ACK Flood to Destination	Excessive number of ACK packets to a destination address.
		RST Flood to Destination	Excessive number of RST packets to a destination address.
		FIN Flood to Destination	Excessive number of FIN packets to a destination address.
		Established Flood to Destination	Excessive number of established connections for a destination address.
		TCP connection	An individual TCP connection has exceeded the packet threshold. Connections are identified using the 4-tuple of Source IP Address, Source TCP Port Number, Destination IP Address, Destination TCP Port Number.
		Legitimate IP	FortiDDoS has determined that an attack is originating from IP addresses it formerly determined to be legitimate. “Zombies” are systems that are unwitting participants in an attack due to infection from a virus or a worm.
		TCP port	Excessive number of packets on a specific TCP port.
		UDP port	Excessive number of packets on a specific UDP port.
		ICMP type/code	Excessive number of ICMP packets.

Table 11: DDoS Attack Log event categories

Layer	Category of event	Traffic parameter	Description
4		Excessive Concurrent Connections Per Source	FortiDDoS has detected excessive instantaneous sources or TCP connections.
		Excessive Concurrent Connections Per Destination	
	ACL	TCP Port	The packet is associated with a TCP or UDP port or ICMP type and code combination that is blocked by the ACL for a Service Protection Profile (SPP).
		UDP Port	
		ICMP type/code	
	Header Anomaly	Layer 4 anomalies	Invalid TCP, UDP, or ICMP checksum Invalid TCP flag combination Urgent flag is set then the urgent pointer must be non-zero SYN, FIN or RST is set for fragmented packets Data offset is less than 5 for a TCP packet EOP (End of packet) is detected before the 20 bytes of TCP header EOP before the data offset indicated data offset Length field in TCP window scale option is a value other than 3 Missing UDP payload Missing ICMP payload
	State Anomaly	Outside TCP/UDP window	The TCP sequence number of a packet was outside the acceptable window.
		Foreign packet	The TCP state machine in FortiDDoS keeps track of all connections through it. A foreign packet is a TCP packet that does not belong to any known connections.
		State transition anomalies	The TCP state machine in FortiDDoS keeps track of all connections through it. An invalid packet is a TCP packet that does not transition correctly from the present state to the next allowed state.
	Scan	Dark Address Scan	TCP blocks packets from any IP-netmask that is denied in the global ACL.

Table 11: DDoS Attack Log event categories

Layer	Category of event	Traffic parameter	Description
	Internal Reason	Excessive hash collisions	An internal event occurs when there are excessive hash collisions, the system is out of memory, or there is some other problem with internal FortiDDoS logic or memory.
		Out of memory	
7	Rate Flood	Method flood	FortiDDoS supports the GET, HEAD, OPTIONS, TRACE, POST, PUT, DELETE and CONNECT methods. Each Service Protection Profile has an incoming and outgoing threshold for each method.
		URL flood	FortiDDoS maintains a hash table of 8 192 URLs for each Service Protection Profile (SPP) in each direction. Each URL in the hash table has corresponding threshold.
		Host, Referer, Cookie, User-Agent flood	FortiDDoS maintains a hash table of 512 Host, Referer, Cookie, and User-Agent header fields per Service Protection Profile in each direction. Each header field value has a corresponding threshold.
		Heuristics flood	HTTP GET requests that do not have the minimum number of HTTP header fields. Source exceeds the maximum number of HTTP accesses that it is allowed to perform during a specified observation period. An IP address exceeds the maximum number of times it is allowed to retrieve the same URL back-to-back (that is, without accessing other URLs in between) within a specified observation period. This behavior is often evidence of a scripting attack, where bots and not humans perform the activity.
		SIP Invite Per Source	Excessive SIP Invite packets from a single source.
		SIP Register Per Source	Excessive SIP Register packets from a single source.
		SIP Concurrent Invite Per Source	Excessive SIP Invite packets from a single source at one time.
	ACL	URL, Host, Referer, Cookie, User-Agent	The packet is associated with a URL or HTTP header field that is blocked by the ACL for a Service Protection Profile (SPP).

Table 11: DDoS Attack Log event categories

Layer	Category of event	Traffic parameter	Description
7	Header Anomaly	Undefined HTTP Method Anomaly	FortiDDoS supports 8 pre-defined methods. FortiDDoS supports 8 pre-defined methods. If an HTTP packet contains method other than these 8 pre-defined methods, it is treated as an unknown HTTP method anomaly.
		Invalid HTTP Version	FortiDDoS supports HTTP versions 1.0 and 1.1. It blocks any packets with any other versions.
NA	Device Events	Disk not OK Hardware Failure Software Failure Disk OK RAID Array Failure RAID Array OK Excessive Events	

See also

- Working with attack reporting
- Viewing log messages
- Analyzing attacks
- Backing up the DDoS attack log
- Deleting DDoS attack log events

Backing up the DDoS attack log

The backup feature allows you to back up DDoS attack log events in the event database by generating a file that you can download.

You can use either the web UI or CLI to generate the events backup file. However, you use the web UI only to download the file.

For more information about the DDoS attack log, see "Logging" on page 221 and "DDoS Attack Log and DDoS Subnet Attack Log" on page 222.

To back up DDoS attack log events via the web UI

1. Click *Log & Report* > *Log Access* > *Log Backup*.
2. For *Service Protection Profiles*, select the appropriate profile.
3. Select *Backup*, and then click *Save*.
4. To confirm that you want to create a new backup file, click *Yes*.
5. After the value of *Status* indicates that the backup file is ready, click *Download*.

   Save or view the file using your browser's download feature.

To back up DDoS attack log events via the CLI

1. Enter the following commands:

```
edit <spp_name>
  config ddos spp attack-event-backup
    set attack-event-backup {enable | disable}
  end
```

   where:

   - `<spp_name>` is the name of the Service Protection Profile (SPP)
   - `{enable | disable}` specifies whether to generate a backup of the DDoS attack log

2. In the web UI, click , and then click *Download* to download the file.

   Save or view the file using your browser's download feature.

See also

- Viewing log messages
- Deleting DDoS attack log events

Deleting DDoS attack log events

After you have backed up the DDoS attack log, you can manually delete log events using a date range. This is useful if you want to make more space available on the hard disk.

You can also configure FortiDDoS to automatically delete older DDoS attack log events when the number of logged events reaches a specified number. This is useful as a safety mechanism.

To delete DDoS attack log events via the web UI

1. Click *Log & Report* > *Log Configuration* > *Purge Settings*.
2. Do one or both of the following:
   - Select *Automatic* and then, for *Purge older events...*, specify a number of events. When the number of recorded events exceeds this value, FortiDDoS deletes older events.
   - Select *Manual*, and then specify start and end dates. FortiDDoS deletes events in this range of dates.
3. Click *Save*.

To delete DDoS attack log events via the CLI

1. Enter the following commands:

```
edit <spp_name>
  config ddos spp attack-event-purge
    [set automatic-event-purge {enable | disable}]
    [set purge-watermark <watermark_int>]
    [set manual-event-purge {enable | disable}]
    [set purge-start-date <purge_date_str>]
    [set purge-end-date <purge_date_str>]
  end
```

   where:

   - `{enable | disable}` specifies whether FortiDDoS purges events automatically or manually
   - `<watermark_int>` is an integer that specifies a number of events. When `automatic-event-purge` is enabled and the number of recorded events exceeds this number, FortiDDoS purges older events.
   - `<purge_date_str>` is a string that specifies the start or end date of the events to purge when `manual-event-purge` is enabled

See also

- Viewing log messages
- Backing up the DDoS attack log

Accessing the DDoS attack log using SQL

You can access the DDoS attack log with read-only permission using a third-party tool such as the MySQL command-line tool or MySQL Workbench.

This feature allows you to view attack log information in a report format other than the one provided by the web UI. For example, to generate consolidated reports when FortiDDoS is integrated with other appliances in your network.

You access the log using the user `root` and the password is the serial number of the appliance.

SQL connections are ***not*** secure, and can be intercepted by a third party. If possible, enable this option only for network interfaces connected to a trusted private network, or directly to your management computer. Failure to restrict administrative access through this protocol could compromise the security of your FortiDDoS appliance.

To enable SQL access to the DDoS attack log via the web UI

1. In the navigation menu, go to *System* > *Network* > *Interface*.
2. Double-click either *mgmt1* or *mgmt2*.
3. For *IPv4/Netmask* or *IPv6/Netmask*, specify the IP address.
4. Under Administrative Access, select *SQL*.

   To allow other types of access to FortiDDoS (for example, HTTPS or SSH), ensure other types of access are selected.

For more information on these settings, see “Configuring the network interfaces” on page 98.

To enable SQL access to the DDoS attack log via the CLI

Enter the following commands:

```
config system interface
  edit {mgmt1|mgmt2}
    set ip <address_ipv4> <netmask_ipv4mask>
    set allow access sql
  next
end
```

where:

- `{mgmt1|mgmt2}` is the interface used to access the DDoS attack log
- `<address_ipv4>` is the IP address assigned to the network interface
- `<netmask_ipv4mask>` is its netmask in dotted decimal format

These commands allow access using SQL only. To allow other types of access to FortiDDoS (for example, HTTPS or SSH), for `set allow access`, specify the additional types of access as required.

For more information on these settings, see “Configuring the network interfaces” on page 98.

To access the DDoS attack log via a command-line interface

The following example illustrates accessing the log using MySQL on a Linux terminal:

```
mysql  -h 172.30.153.122 -u root -pEnter password:
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1393
Server version: 5.5.23-MariaDB Source distribution

...
mysql> show databases;
+--------------------+
| Database           |
+--------------------+
| information_schema |
| flg                |
+--------------------+
2 rows in set (0.00 sec)

mysql> use flg

Reading table information for completion of table and column names
You can turn off this feature to get a quicker startup with -A

Database changed

mysql> select timestamp, inet_ntoa(ip_src4), dropcount from dlog where
dropcount > 1000 order by timestamp desc limit 10;

+---------------------+--------------------+-----------+
| timestamp           | inet_ntoa(ip_src4) | dropcount |
+---------------------+--------------------+-----------+
| 2014-03-13 10:03:07 | 12.0.0.2           |      7471 |
| 2014-03-13 09:47:31 | 10.0.0.2           |      3571 |
| 2014-03-13 09:40:35 | 12.0.0.2           |      3991 |
| 2014-03-13 09:07:29 | 12.0.0.2           |      5649 |
| 2014-03-13 08:38:19 | 10.0.0.2           |      7557 |
| 2014-03-13 07:38:49 | 10.0.0.2           |      2418 |
| 2014-03-13 06:57:48 | 12.0.0.2           |      3425 |
| 2014-03-13 06:57:25 | 10.0.0.2           |      3610 |
| 2014-03-13 06:46:00 | 10.0.0.2           |      1051 |
| 2014-03-13 06:39:12 | 10.0.0.2           |      4853 |
+---------------------+--------------------+-----------+
10 rows in set (0.00 sec)
```

To access the DDoS attack log via graphical user interface

The following example uses MySQL Workbench. You can download MYSQL Workbench for Windows from the following location:

http://dev.mysql.com/downloads/tools/workbench/

1. Open the workbench and log in to the IP address of the appropriate management network interface.

   Use the user `root` and the password is the serial number of the appliance.
2. Open a connection to start querying.
3. In the SQL Editor, select database flg and table dlog.

   For example, query the database using the following query:

   ```
   select dropcount from dlog where dropcount>10000 order by dropcount desc
   ```

System event logs & logging

FortiDDoS appliances log system-related events, including system restarts and HA activity. This information can be useful during troubleshooting or forensic investigation.

You can select a priority level that system event log messages must meet in order to be recorded. For more information, see "System event log severity levels" on page 234.

The FortiDDoS appliance can save system event log messages to its memory, to a remote location in the form of a Syslog server, or both its memory and a remote location. For more information, see "Configuring system event logging" on page 235. The FortiDDoS appliance can also use log messages as the basis for reports. For more information, see "Reports" on page 250.

FortiDDoS appliances can display system event log messages on the dashboard. For more information, see "Event Log Console widget" on page 200.

See also

- System event log severity levels
- Configuring system event logging
- Selecting which system events to log
- Configuring logging to a remote logging server
- Viewing log messages

System event log severity levels

Each system event log message contains a *Priority* (`pri`) field that indicates the severity of the event that caused the log message, such as `pri=warning`.

Table 12: System event log severity levels

Level (0 is greatest)	Name	Description
0	Emergency	The system has become unusable.
1	Alert	Immediate action is required.
2	Critical	Functionality is affected.
3	Error	An error condition exists and functionality could be affected.
4	Warning	Functionality could be affected.

Table 12: System event log severity levels

Level (0 is greatest)	Name	Description
5	Notification	Information about normal events.
6	Information	General information about system operations.

For each location where the FortiDDoS appliance can store system event log files (disk, memory, Syslog), you can define a severity threshold. The FortiDDoS appliance stores all log messages equal to or exceeding the log severity level that you select.

For example, if you select *Error*, the FortiDDoS appliance stores log messages with a log severity level of *Error*, *Critical*, *Alert*, and *Emergency*.

Avoid recording log messages using low log severity thresholds (such as information or notification) to the local hard disk for an extended period of time. A low log severity threshold is one possible cause of frequent logging. Excessive logging frequency can cause undue wear on the hard disk and may cause premature failure.

For more information, see "Configuring logging to a remote logging server" on page 237.

See also

- Configuring system event logging
- Selecting which system events to log
- Configuring logging to a remote logging server
- Viewing log messages

Configuring system event logging

You can configure the FortiDDoS appliance to store system event log messages either locally (that is, in RAM or to the hard disk), remotely (that is, on a Syslog server), or both locally and remotely. Several factors determine your choice of storage location, including the following factors:

- Logging only locally may not satisfy your requirements for off-site log storage.
- Very frequent logging on the local hard drive can cause undue wear on the drive. A low severity threshold is one possible cause of frequent logging. For more information on severity levels, see "System event log severity levels" on page 234.
- Very frequent logging, such as when the severity level is low, can rapidly consume all the memory that is available for logs. If the available space is consumed, and if the FortiDDoS appliance is configured to do so, it may store any new log message by overwriting the oldest log message. When traffic volumes are high, this overwriting process can occur so rapidly that you cannot view old log messages before they are replaced.
- In most cases, you can reduce the number of log messages that are stored in memory. Logging to a Syslog server can provide you with additional log storage space.

For information on viewing locally stored log messages, see "Viewing log messages" on page 239.

To configure logging

1. If you store logs remotely, configure connectivity information such as the IP address. See “Configuring logging to a remote logging server” on page 237.
2. Specify the types of system events that you want to record to the specified local or remote destinations. See “Selecting which system events to log” on page 236.
3. Use the web UI or alert mail to monitor your log messages for events that require action from network administrators. See “Viewing log messages” on page 239 and “Alert email” on page 241. Configure reports that are derived from log data to review trends in your network. See “Reports” on page 250.

You cannot use the FortiDDoS web UI to view system event logs that are stored remotely. If you require the ability to view logs from the web UI, also enable local storage. For details, see “Selecting which system events to log” on page 236.

See also

- System event log severity levels
- Selecting which system events to log
- Configuring logging to a remote logging server
- Viewing log messages

Selecting which system events to log

Log & Report > *Log Configuration* > *Log Settings* allows you to specify which system events to log or to disable system event logging completely.

To select system events to log

1. Go to *Log & Report* > *Log Configuration* > *Log Settings*.

   To access this part of the web UI, your administrator’s account access profile must have *Read-Write* permission to items in the *Log & Report* category. For details, see “Restricting permissions” on page 174.
2. Configure the types and severity of system event log messages that FortiDDoS records.

3. To disable system event logging, clear the *Event Logging* option.

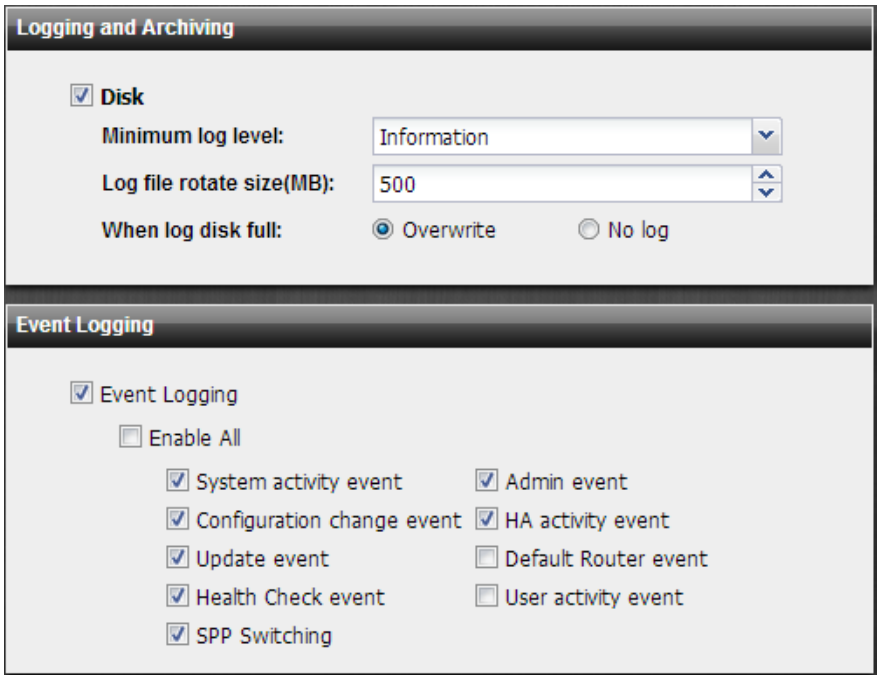

4. Click *Save*.
5. Logs are recorded to the hard disk of the appliance. If you also want logs to be recorded to a Syslog server, continue with "Configuring logging to a remote logging server" on page 237.

See also

- System event log severity levels
- Configuring system event logging
- Configuring logging to a remote logging server
- Viewing log messages

Configuring logging to a remote logging server

To store system event and DDoS attack log event logs in a safe remote location or offload logging for performance reasons, you can configure FortiDDoS to store logs on a generic Syslog server.

To ensure logging accuracy, verify that the FortiDDoS appliance's system time is accurate. For details, see "Setting the system time & date" on page 95.

Avoid recording highly frequent log types such as traffic logs to the local hard disk for an extended period of time. Excessive logging frequency can cause undue wear on the hard disk and premature failure.

To configure system event logging to a remote location

1. Go to *Log & Report* > *Log Configuration* > *Log Remote*.

   To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *Log & Report* category. For details, see "Restricting permissions" on page 174.

2. Click *Add*.
3. Select *Enable*.
4. In *Address*, type the address of the Syslog server.
5. In *Port*, type the UDP port number where the device listens for logs. The default is 514, which is the typical default value for Syslog.
6. If the Syslog server supports logging in comma-separated value (spreadsheet) format, you can enable *CSV Format*.
7. Select the *Minimum Log Level* and system event log types that a message must match in order to be sent to the remote server.
8. From *Facility*, select an identifier that is not used by any other device on your network when sending logs to Syslog.
9. Click *Save*.
10. To verify logging connectivity, from the FortiDDoS appliance, trigger a log message that matches the types and severity levels that you have chosen to store on the remote host. Then, on the remote host, confirm that it has received that log message.

    If the remote host does not receive the log messages, verify the FortiDDoS appliance's network interfaces (see "Configuring network interfaces, gateway, and DNS" on page 98) and static routes (see "Adding a gateway" on page 103), and the policies on any intermediary firewalls or routers. If ICMP `ECHO_RESPONSE` (pong) is enabled on the remote host, try using the `execute traceroute` command to determine the point where connectivity fails.

To configure DDoS attack log event logging to a remote location

1. Go to *Log & Report* > *Log Configuration* > *DDoS Attack Log Remote*.

   To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *Log & Report* category. For details, see "Restricting permissions" on page 174.
2. Click *Add*.
3. For *Name*, enter the name that identifies this service protection profile (SPP) and logging server combination in the *DDoS Attack Log Remote Configuration* list.
4. Select *Enable*.
5. For *SPP*, select the profile whose logs are stored in the remote location.

   You can specify only one remote logging destination for each SPP.
6. In *Address*, type the address of the remote Syslog server.
7. For *Port*, enter the UDP port number where the device listens for logs. The default value is 514, which is the typical default value for Syslog.
8. Click *Save*.
9. If you specified the address of remote Syslog server, confirm with the server administrator that the FortiDDoS appliance was added to the server's device list, allocated sufficient disk space quota, and assigned permission to transmit logs to the server.
10. To verify logging connectivity, generate a rate flood event for the specified profile. Then, on the remote host, confirm that it has received the corresponding DDoS Attack Log message.

    If the remote host does not receive the log messages, verify the FortiDDoS appliance's network interfaces (see "Configuring network interfaces, gateway, and DNS" on page 98) and static routes (see "Adding a gateway" on page 103), and the policies on any intermediary firewalls or routers. If ICMP `ECHO_RESPONSE` (pong) is enabled on the remote host, try using the `execute traceroute` command to determine the point where connectivity fails.

See also

- System event log severity levels
- Configuring system event logging
- Selecting which system events to log
- Viewing log messages

Viewing log messages

You can use the web UI to view and download locally stored log messages. (You cannot use the web UI to view log messages that are stored remotely on a Syslog server.)

To view log messages

Go to one of the logs:

- *Log & Report* > *Log Access* > *Event Log*
- *Log & Report* > *Log Access* > *DDoS Attack Log*

To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *Log & Report* category. For details, see "Restricting permissions" on page 174.

Columns and appearance varies by log type.

Initially, the page displays the most recent log messages for that log type.

In FortiDDoS HA clusters, log messages are recorded on their originating appliance. If you notice a gap in the logs, a failover may have occurred. Logs for that period are stored on the other appliance. To view those logs, switch to the other appliance.

See also

- System event logs & logging
- Displaying & arranging log columns
- Filtering log messages

Displaying & arranging log columns

When viewing logs, you can show, hide and re-order most columns to display only relevant categories of information in your preferred order.

For most columns, you can also filter data within the columns to include or exclude log messages which contain your specified text in that column. For more information, see "Filtering log messages" on page 240.

To rearrange columns, click and drag the column headers into the order that you prefer.

To display or hide columns

1. Go to one of the log types, such as *Log & Report* > *Log Access* > *Event Log*.

   To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *Log & Report* category. For details, see "Restricting permissions" on page 174.
2. Click the arrow on the right side of a column header.

3. In the drop-down menu that appears, select *Columns* option.
4. In the list of available columns that appears, enable (show) or disable (hide) each column.

   The page refreshes, displaying the columns that you selected, in the order that you specified. Column settings persist when changing pages or logging out, and apply to all administrator accounts with access to the page.

See also

- Filtering log messages

Filtering log messages

You can filter columns to display only those log messages that do or do not contain your specified content in that column.

To filter log messages by column contents

1. Go to one of the log types, such as *Log & Report* > *Log Access* > *Event Log*.

   To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *Log & Report* category. For details, see "Restricting permissions" on page 174.
2. On the tool bar, click *Filter Settings*.

   The filter dialog appears.
3. Click the green + (plus) icon next to *Add New Filter*.
4. From the *Field* list, select the name of the column to use to filter the log view.
5. Do the following as required:
   - To ***exclude*** log messages with matching content in this column, select *Exclude*.
   - For *Date* and *Time* filters, define the time period in *To* and *From*.
   - For other filters, in *Value*, select or type the ***entire*** value that matching log messages must contain in that column. Appropriate filter strings vary by the column.

Type the ***entire*** value of a field in the column ***exactly***, or use wild card characters (*) to indicate multiple possible matching values. Otherwise, either results will be different than you intend, or no log messages may entirely match the filter, and so the results will be empty.

For example, when filtering the log view based upon the *ID* column, in *Value*, you can type an entire, single log ID:

```
00032009
```

or you could match multiple log IDs by using an asterisk (*) to match multiple characters:

```
*32009
```

```
00032*
```

```
*32*
```

Matching log messages are excluded or included in your view based upon whether you have selected *Exclude*.

6. Click *Apply* to create the filter.

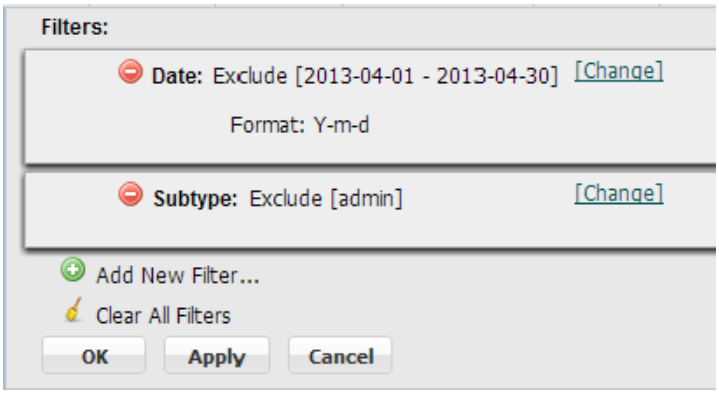

7. Click *OK* when you have added all the filters you require.

To clear one filter

1. Go to one of the log types, such as *Log & Report* > *Log Access* > *Event Log*.
2. Click the *Filter Settings* button.

   The filter dialog appears.
3. Click the red minus sign (-) next to the column name in the filter dialog.
4. Click *OK*.

   The page refreshes. It should now include logs that were previously filtered.

To clear all filters

1. Go to one of the log types, such as *Log & Report* > *Log Access* > *Event Log*.
2. On the tool bar, click the *Filter Settings* button.
3. Click *Clear All Filters*.
4. Click *OK*.

   The page refreshes, and should now display all logs of that type, unfiltered.

See also

- Displaying & arranging log columns

Alert email

To notify you of serious system failure events, you can configure the FortiDDoS appliance to generate an alert email.

System event alerts appear on the dashboard. FortiDDoS can also generate alert email if you configure email settings and include them in a trigger that is used by system resource thresholds.

System events that generate alert email also generate log messages. If you have received an alert email and want to know more about the events, go to the corresponding log messages. For information on viewing locally stored log messages, see "Viewing log messages" on page 239.

The alert email settings also specify the destination for the email output for reports. For more information, see "Selecting the report's file type & email delivery" on page 257.

To configure alert email

1. Go to *Log & Report* > *Log Configuration* > *Alert Mail*.
2. Configure the following settings:

SMTP server	Enter the SMTP server address.
SMTP port	Enter the SMTP server port number. The default value is port 25.
Email from	Enter an email address. This address appears in the from field in the email message.
Email to	Enter up to three recipient email addresses.
Authentication	Select if authentication is required.
SMTP user	Enter the SMTP username.
Password	Enter the SMTP password.
Test Connectivity	Select to test connectivity to the configured SMTP server.

3. Click *Save*.
4. Go to *Log & Report* > *Log Configuration* > *Alert Mail*.
5. Select one of the following options:
   - *Send alert email for the following* (to send alert email for specific network events)
   - *Send alert email for logs based on severity*
6. Configure the following settings as required:

Interval Time	Enter the interval time. The default value is 30 minutes.
Send alert email for the following	You can select to configure alert email for the following network events: • HA status changes • Administrator login/logout • Configuration changes • Diskfull • Health check • Update • Default gateway changes • User changes • SPP Switching
Send alert email for logs based on severity	Select the minimum log level from the drop down menu. See "System event log severity levels" on page 234.

7. Click *Save*.

See also

- Configuring system event logging
- Selecting the report's file type & email delivery

SNMP traps & queries

The FortiDDoS SNMP settings allow you to configure the FortiDDoS to allow queries for information and to send traps to one or more computers that you specify as SNMP managers. In this way you can use an SNMP manager to monitor the FortiDDoS appliance.

FortiDDoS can transmit two types of simple network management protocol (SNMP) messages (traps) to an SNMP manager. You configure the SNMP settings for each type of message separately:

- System alarms and event messsages
- Attack log messages

Before you can use SNMP, you must configure the FortiDDoS appliance's SNMP settings and add it to at least one community. An SNMP community is a grouping of equipment for network monitoring purposes. Your FortiDDoS appliance has to be a member of least one SNMP community so that community's SNMP managers can query the FortiDDoS appliance's system information and receive SNMP traps from the FortiDDoS appliance.

You must also enable SNMP access on the network interface that the SNMP manager uses to connect to FortiDDoS. (See "Configuring network interfaces, gateway, and DNS" on page 98.)

For system messages, you use *System* > *Config* > *SNMP* to specify a SNMP manager and other SNMP settings.

To specify machines that receive FortiDDoS attack log information via SNMP, use *Log & Report* > *Log Configuration* > *SNMP Trap Receivers* . You can specify one or more SNMP managers for the attack logs from each service protection profile (SPP). FortiDDoS can send the attack log information for an SPP to more than one manager.

On the SNMP manager, you ensure that the SNMP manager is a member of the community that the FortiDDoS appliance belongs to, and compile the necessary Fortinet-proprietary management information blocks (MIBs) and Fortinet-supported standard MIBs. For information on MIBs, see "MIB support" on page 249.

If you do not configure the SNMP manager as a host in a community to which the FortiDDoS appliance belongs, or to supply it with required MIBs, the SNMP monitor is unable to query or receive traps from the FortiDDoS appliance.

Configuring SNMP settings for system alarms and event messages

To configure the SNMP settings for system messages

1. Add the MIBs to your SNMP manager so that you can receive traps and perform queries. For instructions, see the documentation for your SNMP manager.
2. Go to *System* > *Config* > *SNMP*.

   To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *System* category. For details, see "Restricting permissions" on page 174.

3. Configure the following:

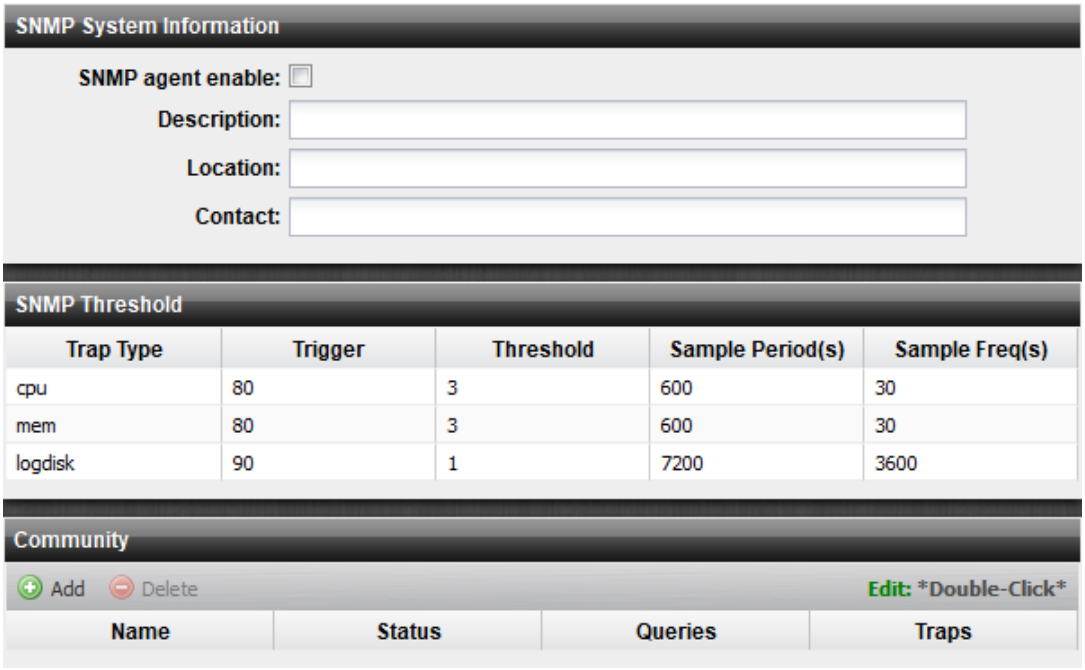

Setting name	Description
SNMP Agent	Enable to activate the SNMP agent, so that the FortiDDoS appliance can send traps and receive queries for the communities in which you enabled queries and traps.
Description	Type a comment about the FortiDDoS appliance, such as `dont-reboot`. The description can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).
Location	Type the physical location of the FortiDDoS appliance, such as `floor2`. The location can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).
Contact	Type the contact information for the administrator or other person responsible for this FortiDDoS appliance, such as a phone number (`555-5555`) or name (`jdoe`). The contact information can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).

4. Click *Save*.
5. Create at least one SNMP community to define which hosts are allowed to query and which hosts receive traps. See "To add a community to the SNMP configuration for system messsages" on page 244.

To add a community to the SNMP configuration for system messsages

The SNMP configuration for system messages can have up to three SNMP communities. Each community can have a different configuration for queries and traps, and the set of events that trigger a trap. You can also add the IP addresses of up to eight SNMP managers to each community to designate the destination of traps and which IP addresses are permitted to query the FortiDDoS appliance.

1. Go to *System* > *Config* > *SNMP*.

   To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *System* category. For details, see “Restricting permissions” on page 174.
2. Ensure the SNMP agent is configured. For more information, see “To configure the SNMP settings for system messages” on page 243.
3. Under Community, click *Add*.

   A dialog appears.

4. Configure these settings:

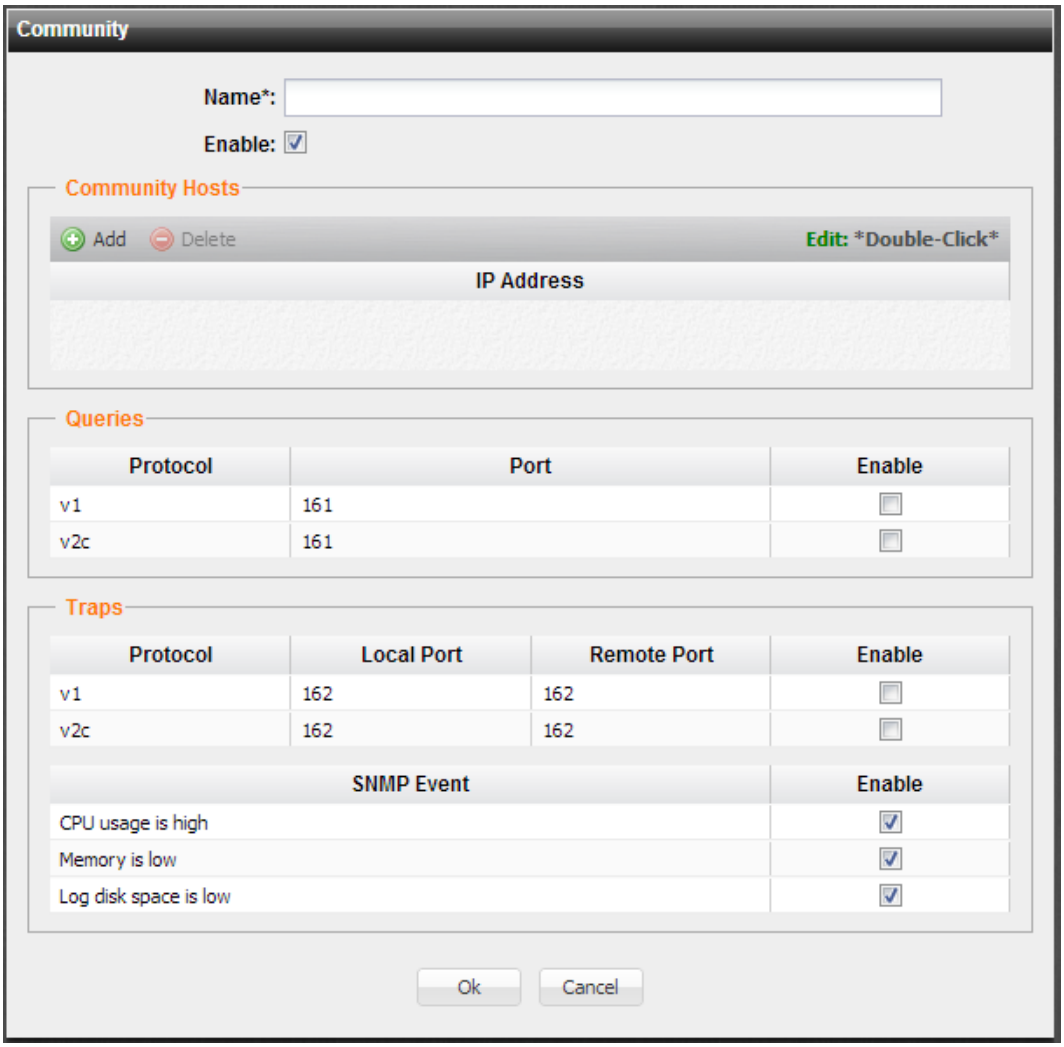

Setting name	Description
Community Name	Type the name of the SNMP community to which the FortiDDoS appliance and at least one SNMP manager belongs (for example, `public`). The FortiDDoS appliance does not respond to SNMP managers whose query packets do not contain a matching community name. Similarly, trap packets from the FortiDDoS appliance include community name, and an SNMP manager may not accept the trap if its community name does not match. **Caution:** Fortinet strongly recommends that you do ***not*** add FortiDDoS to the community named `public`. This popular default name is well-known, and attackers that gain access to your network often try this name first.

Setting name	Description
Hosts	
IP Address	Click *Add* and enter the IP address of the SNMP manager that, if traps or queries are enabled in this community: • receives traps from the FortiDDoS appliance • is permitted to query the FortiDDoS appliance SNMP managers have read-only access. To allow any IP address using this SNMP community name to query the FortiDDoS appliance, enter `0.0.0.0`. For security best practice reasons, however, this is not recommended. **Caution:** FortiDDoS sends security-sensitive traps, which should be sent only over a trusted network and only to administrative equipment. **Note:** If there are no other host IP entries, entering only `0.0.0.0` effectively disables traps because there is no specific destination for trap packets. ***If you do not want to disable traps, you must add at least one other entry*** that specifies the IP address of an SNMP manager.
Delete	Select an IP address and click *Delete* to remove an SNMP manager from the SNMP community configuration.
Add	Click to add additional SNMP manager entries. You can add up to 8 SNMP managers to each community.
Queries	To edit the port number on which the FortiDDoS appliance listens for SNMP queries from the SNMP managers in this community, click the current value (161 by default). Then, enable queries for either or both SNMP v1 and SNMP v2c.
Traps	To edit the port number that is the source (*Local*) port number and destination (*Remote*) port number for trap packets sent to SNMP managers in this community, click the current value (162 by default). Then, enable traps for either or both SNMP v1 and SNMP v2c.
SNMP Event	Enable the types of SNMP traps that you want the FortiDDoS appliance to send to the SNMP managers in this community. While most trap events are described by their names, the following events occur when a threshold has been exceeded: • **CPU usage is high** — CPU usage has exceeded 80%. • **Memory is low** — Memory (RAM) usage has exceeded 80%. • **Log disk space is low** — Disk space usage for the log partition or disk has exceeded 90%. For more information on supported traps and queries, see "MIB support" on page 249.

5. Click *OK*.
6. To verify your SNMP configuration and network connectivity between your SNMP manager and your FortiDDoS appliance, be sure to test both traps and queries (assuming you have enabled both).

   Traps and queries typically occur on different port numbers, and therefore verifying one does not necessarily verify that the other is also functional. To test queries, from your SNMP manager, query the FortiDDoS appliance. To test traps, cause one of the events that should trigger a trap.

See also

- Configuring network interfaces, gateway, and DNS
- MIB support
- Configuring SNMP settings for attack log messages

Configuring SNMP settings for attack log messages

1. Add the MIBs to your SNMP manager so that you can receive traps. For instructions, see the documentation for your SNMP manager.
2. *Go to Log & Report > Log Configuration > SNMP Trap Receivers*.

   To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *Log & Report* category. For details, see "Restricting permissions" on page 174.
3. Click *Add* and complete the following settings:

Name		Enter a name that identifies this SNMP trap receiver in the list of receivers.
Enable		Specify whether this SNMP configuration is enabled.
SPP		Specify the service protection profile that is the source of attack log messages that FortiDDoS sends to the SNMP manager.
Address		Enter the IP addresss of the SNMP manager that receives attack log traps.
Port		Specify the port on which the SNMP manager listens for attack log information. The default value is 162.
Community Username		Type the name of the SNMP community to which the FortiDDoS appliance and the SNMP manager at the specified address belong.
SNMP Version		Specifies the version of SNMP used for attack log traps.
	Engine Id	If *SNMP Version* is *v3*, specifies the engine ID that uniquely identifies the SNMP agent.
	v3 Access Type	Specifies whether the SNMP agent communicates with authentication.

4. Click *Save*.
5. Add additional entries as required.

   For example, you can add additional SNMP managers that receive traps from the same SPP or additional managers that receive traps from different SPPs.

See also

- MIB support
- Configuring SNMP settings for system alarms and event messages

MIB support

The FortiDDoS SNMP agent supports a few management information blocks (MIBs).

Table 13: Supported MIBs

MIB or RFC	Description
Fortinet Core MIB	This Fortinet-proprietary MIB enables your SNMP manager to query for system information and to receive traps that are common to multiple Fortinet devices.
FortiDDoS MIB	This Fortinet-proprietary MIB enables your SNMP manager to query for FortiDDoS-specific information and to receive FortiDDoS-specific traps.
RFC 1213 (MIB II)	The FortiDDoS SNMP agent supports MIB II groups, except: • There is no support for the EGP group from MIB II (RFC 1213, section 3.11 and 6.10). • Protocol statistics returned for MIB II groups (IP, ICMP, TCP, UDP, and so on) do not accurately capture all FortiDDoS traffic activity. More accurate information can be obtained from the information reported by the FortiDDoS MIB.
RFC 2665 (Ethernet-like MIB)	The FortiDDoS SNMP agent supports Ethernet-like MIB information, except the dot3Tests and dot3Errors groups.
RFC 3411	SNMP v4 Architecture for SNMP Frameworks
RFC 3414	Partial support for User-based Security Model

You can obtain these MIB files from the Fortinet Technical Support web site, https://support.fortinet.com/.

To communicate with your FortiDDoS appliance's SNMP agent, you must first compile these MIBs into your SNMP manager. If the standard MIBs used by the SNMP agent are already compiled into your SNMP manager, you do not have to compile them again. The FortiDDoS SNMP implementation is read-only.

To view a trap or query's name, object identifier (OID), and description, open its MIB file in a plain text editor.

All traps sent include the message, the FortiDDoS appliance's serial number, and host name.

For instructions on how to configure traps and queries, see "SNMP traps & queries" on page 243.

See also

- SNMP traps & queries

Reports

FortiDDoS generates report information using log messages that it has recorded. It displays the information either in a report format that you configure or widgets on the *Executive Summary* (*Log & Report* > *Report Browse* > *Executive Summary*) or *Subnet Executive Summary* (*Log & Report* > *Report Browse* > *Executive Summary*) dashboards.

FortiDDoS can generate reports automatically according to the schedule that you create in the report configuration. You can also generate them manually by clicking the *Run now* icon in the report profile list.

Generating reports can be resource intensive. To avoid problems with traffic processing performance, generate reports during times with low traffic volume, such as at night or weekends. For more information on scheduling the generation of reports, see "Scheduling reports" on page 257.

Viewing report information on a dashboard (Executive Summary)

FortiDDoS has two dashboards that allow you to quickly view DDoS attack activity report information in a table format: *Executive Summary* and *Subnet Executive Summary*.

To see the report information as a set of dasboard widgets, click *Log & Report* > *Report Browse* > *Executive Summary*.

In the *Executive Summary* dashboard, the widgets display the top 30 attacks in each report category.

To see all attacks in the top attacked subnet and top ACL subnet drops category, click *Log & Report* > *Report Browse* > *Subnet Executive Summary*.

For information on managing the dashboard widgets, see "The dashboard" on page 195.

For a description of the information in each widget, see "DDoS Attack Activity report types" on page 255.

Configuring a report

A FortiDDoS report configuration specifies the settings that the appliance should use, such as the report name, file format, and subject matter.

To configure a report

1. For a system event report, collect the log data that the report is based on. For information on enabling logging to the local hard disk, see "Configuring system event logging" on page 235.

   For DDoS attack event reports, FortiDDoS generates the log data automatically.
2. Go to *Log & Report* > *Report Configuration* > *Report Configuration*.

   To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *Log & Report* category. For details, see "Restricting permissions" on page 174.

3. Click *Add*.

   A dialog appears.

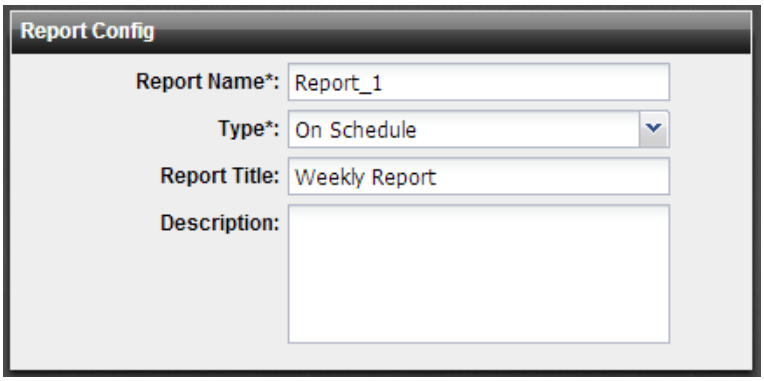

4. In *Report Name*, type a name for the report. The name cannot contain spaces.
5. If you are creating a report profile, select from *Type* either to run the report immediately after configuration (*On Demand*) or run the report at configured intervals (*On Schedule*). This cannot be changed later.

For on-demand reports, the FortiDDoS appliance does ***not*** save the report profile after it generates the report. If you want to save the report profile, but do not want to generate the report at regular intervals, select *On Schedule*, but then in the *Schedule* section, select *Not Scheduled*.

6. In *Report Title*, type the name that appears in the title area of the report. The title can include spaces.
7. In *Description*, type a comment or other description.
8. Configure the following report options as required:

Setting name	Description
Properties	Adds custom logos, headers, footers and company information to the report. For more information, see "Customizing the report's headers, footers, & logo" on page 252.
Report Scope	Specifies the time period covered by the report.
Report Type(s)	Specifies the types of system event messages to include in the report. For more information, see "Choosing the type & format of a report profile" on page 254.
DDoS Attack Activity	Specifies the types of DDoS attack events to include in the report. For more information, see "DDoS Attack Activity report types" on page 255.
Report Format	Specifies the number of top items to include in ranked report subtypes, and other advanced features. For more information, see "Choosing the type & format of a report profile" on page 254.
Schedule	Specifies when the FortiDDoS appliance runs the report. For example, daily or on specific days of the week. For more information, see "Scheduling reports" on page 257.
Output	Specifies the file formats for reports generated by this report profile. For more information, see "Selecting the report's file type & email delivery" on page 257.

9. Click *Save*.

   The appliance generates on-demand reports immediately.

   Scheduled reports are generated at intervals set in the schedule. For information on viewing generated reports, see "Viewing & downloading generated reports" on page 258.

To generate a report immediately

1. In the *Report Config* list, select the check box for the report.
2. Click *Run now*.

See also

- Customizing the report's headers, footers, & logo
- Restricting the report's scope
- Choosing the type & format of a report profile
- Scheduling reports
- Selecting the report's file type & email delivery

Customizing the report's headers, footers, & logo

When configuring a report profile, you can provide text and logos to customize the appearance of reports generated from the profile.

Setting name	Description
Company Name	Type the name of your company or other organization.
Header Comment	Type a title or other information to include in the header.
Footer Comment	Select which information to include in the footer: • *Report Title* — Use the text from *Report Name*. • *Custom* — Use other text that you type into the field to the right of this option.
Title Page Logo	Select *No Logo* to omit the title page logo. Select *Custom* to include a logo, and then click *Select* to locate the logo file, and click *Upload* to save it to the FortiDDoS appliance's hard disk for use in the report title page. See "To upload a logo file".
Header Logo	Select *No Logo* to omit the header logo. Select *Custom* to include a logo, and then click *Select* to locate the logo file, and click *Upload* to save it to the FortiDDoS appliance's hard disk for use in the report header. The header logo will appear on every page in PDF- and Microsoft Word (RTF)-formatted reports, and at the top of the page in HTML-formatted reports.

To upload a logo file

1. In the *Properties* section of the *Log Report Config* dialog, for either *Title Page Logo* or *Header Logo*, select *Custom*.

2. Click *Select*.

   A dialog appears.

3. To select the logo file on your computer, click *Choose File*.
4. Click *Upload*.

   A rendering of the logo appears in the dialog.

5. Select the logo and click *OK*.

   The name of the logo appears next to *Custom* on the *Log Report Config.*

When adding a logo to the report, select a logo file format that is compatible with your selected file format outputs. If you select a logo that is not supported for a file format, the logo will not appear in that output. For example, if you provide a logo graphic in WMF format, it will not appear in PDF or HTML output.

Table 14: Report file formats and their supported logo file formats

PDF reports	JPG, PNG, GIF
RTF reports	JPG, PNG, GIF, WMF
HTML reports	JPG, PNG, GIF

To delete a logo file

1. Expand the *Properties* section of the *Log Report Config* dialog.
2. For either *Title Page Logo* or *Header Logo*, click the *Select* link beside the logo name you want to remove.

   A dialog appears.

3. Select the logo to remove.
4. Click *Delete*.

Restricting the report's scope

When you configure a report profile, you can select the time span of log messages from which to generate the report. (To start at the beginning of the report configuration instructions, see "Configuring a report" on page 250.)

Setting name	Description
Time Period	Select one of the following: • **Not Used**: Includes all events, regardless of their date • **Absolute**: Allows you to specify a date and time range to include. • **A time span:** (for example, *Today*, *This Month* or *Past N Days*.
Past *N* Hours **Past *N* Days** **Past *N* Weeks**	Enter the number ***N*** of the unit of time. This option appears only when you have selected *Last N Hours*, *Last N Days*, or *Last N Weeks* from *Time Period*, and therefore must define ***N***.

Setting name		Description
	From Date Hour	Select and configure the beginning of the time span. For example, you may want the report to include log messages starting from May 5, 2006 at 6 PM. You must also configure *To Date*.
	To Date Hour	Select to configure the end of the time span. For example, you may want the report to include log messages up to May 6, at 12 AM. You must also select and configure *From Date*.

Choosing the type & format of a report profile

When you configure a report profile, you can:

- Select one or more queries or query groups that define the subject matter of the report
- Configure various advanced options that affect how many log messages are used to formulate ranked report subtypes
- Configure how results are displayed

(To start at the beginning of the report configuration instructions, see "Configuring a report" on page 250.)

Setting name	Description
Report Type(s)	Each of the options in the Event Activity and DDoS Attack Activity query groups corresponds to a chart that appears in the generated report. To select all queries within the group, select the query group. You can also select queries individually. For example: • To include charts about both normal traffic and attacks, select *Enable All* under both *Event Activity* and *DDoS Attack Activity*. • To include only a chart about top system event types only, under *Event Activity*, select *Top Event Types* only. For the DDoS Attack Activity queries, you also specify whether the report is generated for all Service Protection Profiles (SPPs) or a specific profile. For a description of the DDoS Attack Activity queries, see "DDoS Attack Activity report types" on page 255
Report Format	

Setting name	Description
In 'Ranked Reports' show top	Ranked reports (top ***x***, or top ***y*** of top ***x***) can include a different number of results per cross-section, and then combine remaining results under "Others." For example, in *Top Sources By Top Destination*, the report includes the top ***x*** destination IP addresses, and their top ***y*** source IP addresses, and then groups the remaining results. You can configure both ***x*** and ***y*** in *Report Format.* In ranked reports, ("top ***x***" report types, such as *Top Attack Type*), you can specify how many items from the top rank are included in the report. For example, you can set the *Top Attack URLs* report to include up to 30 of the top ***x*** denied URLs by entering `30` for *values of the first variable 1.. 30*. Some ranked reports rank two aspects. For example, the *Top Sources By Top Destination* report ranks top source IP addresses for each of the top destination IP addresses. For these double-ranked reports, you can also configure the rank threshold of the second aspect by entering the second threshold in *values of the second variable for each value of the first variable 1..30*. **Note:** Reports that do not include "Top" in their name display all results. Changing the ranked reports values does not affect these reports.
values of the first variable 1.. 30	Type the value of ***x***.
values of the second variable for each value of the first variable 1.. 30	Type the value of ***y***. This value is only considered if the report rankings are nested (that is, top ***y*** of top ***x***).
Include Summary Information	Enable to include a listing of the report profile settings.
Include Table of Contents	Enable to include a table of contents for the report.

DDoS Attack Activity report types

The following DDoS Attack Activity options allow you to filter the report information by service protection profile (SPP) or subnet.

- **All SPPs** — Report includes information for all profiles.
- **SPP** — Report includes information for the profile you specify only.
- **Subnet** — Report includes information for the subnet you specify only.
- **Default Subnet** — Report includes information for the default subnet only. The default subnet is the subnet that is defined in the SPP policy that has a Subnet ID value of 0.

For information on defining subnets, see "Create a service protection profile (SPP)" on page 113.

When you filter information by subnet, only the *Top Attacks* and *Top ACL Attacks* report types are available:

- Top Attacks

  This report depicts network attacks according to the type of attack in descending order of packets dropped.

- Top ACL Attacks

  This report depicts traffic blocked by ACL configuration according to the type of attack in descending order of packets dropped.

- Top Attackers

  This report depicts network attacks where the source IP address has been tracked. It is shown in descending order of packets dropped.

- Top Attacked Destinations

  This report depicts network attacks where the destination IP address has been tracked. It is shown in descending order of packets dropped.

- Top Attacked Subnet/Top Attacks

  This report depicts network attacks that occurred on a specific subnet. It is shown in descending order of packets dropped. The report identifies subnets using the subnet ID that you assigned in the *SPP Policy* settings.

- Top Attacked Protocols

  This report depicts network attacks that occurred on particular IP protocols in descending order of packets dropped.

- Top Attacked TCP Ports

  This report depicts network attacks that occurred on particular TCP ports in descending order of packets dropped.

- Top Attacked UDP Ports

  This report depicts network attacks that occurred on particular UDP ports in descending order of packets dropped.

- Top Attacked ICMP Type Codes

  This report depicts network attacks that occurred on particular ICMP Types and Codes in descending order of packets dropped.

- Top Attacked HTTP Methods

  This report depicts network attacks where the HTTP method has been identified. It is shown in descending order of packets dropped.

- Top Attacked HTTP URLs

  This report depicts attacks on a specific URL. It is shown in descending order of packets dropped.

- Top Attacked HTTP Hosts

  This report depicts network attacks where FortiDDoS has identified the Host header field. It is shown in descending order of packets dropped.

- Top Attacked HTTP Referrers

  This report depicts network attacks where FortiDDoS has identified the Referer header field. It is shown in descending order of packets dropped.

- Top Attacked HTTP Cookies

  This report depicts network attacks where FortiDDoS has identified the Cookie header field. It is shown in descending order of packets dropped.

- Top Attacked HTTP User Agents

This report depicts network attacks where FortiDDoS has identified the User-Agent header field. It is shown in descending order of packets dropped.

Most of the DDoS Attack Activity report information is also available as a graph (*System* > *Status* > *SPP Attacks*) and in tables provided by the *Executive Summary* dand Subnet Executive Summary dashboards (*Log & Report* > *Report Browse* > *Executive Summary*/*Subnet Executive Summary*).

Scheduling reports

When you configure a report profile, you can select whether the FortiDDoS appliance generates the report on demand or according to the schedule that you configure. (To start at the beginning of the report configuration instructions, see “” on page 250.)

Generating reports can be resource-intensive. To improve performance, schedule reports during times when traffic volume is low, such as at night or during weekends. To determine the current traffic volume, on the dashboard, see the *System Resources* widget.

Setting name		Description
Schedule		
	Not Scheduled	Select if you do ***not*** want the FortiDDoS appliance to generate the report automatically according to a schedule. If you select this option, you can only generate the report by clicking *Run now* in report profile list. For more information, see “Reports” on page 250.
	Daily	Select to generate the report each day. Also configure *Time*.
	These Days	Select to generate the report on specific days of each week, and then mark the check boxes for those days. Also configure *Time*.
	These Dates	Select to generate the report on specific date of each month, and then enter those date numbers. Separate multiple date numbers with a comma. Also configure *Time*. For example, to generate a report on the first and 30th day of every month, enter `1,30`.
Time		Select the time of the day when the report is generated. This option does not apply if you have selected *Not Scheduled*.

Selecting the report’s file type & email delivery

When you configure a report profile, you can select one or more file formats in which to save reports generated by the profile. You can also configure the FortiDDoS appliance to email the report files to specific recipients. (To start at the beginning the report configuration instructions, see “” on page 250.)

To enable the email functionality, ensure that any values that you select for Email Output are also selected for File Output.

Setting name	Description
File Output	Specify which file formats that you want to generate and store on the FortiDDoS appliance's hard drive. Because FortiDDoS always generates the HTML reports, the HTML option is always selected. You can also select one or more of the following additional file formats: • plain text (*Text*), and • *PDF* • *MS Word* (RTF) • MIME HTML (*MHT*, which can be included in email)
Email Output	Select the file formats that you want to generate for an email that is sent to the recipients specified by *Log & Report* > *Log Configuration* > *Alert Mail*. Make sure that the values you specify are also selected for *File Output*.

Viewing & downloading generated reports

Log & Report > *Report Browse* > *Report Browse* displays a list of generated reports that you can view, delete, and download.

In FortiDDoS HA clusters, generated reports (PDFs, HTML, RTFs, plain text, or MHT) are recorded on their originating appliance. If you cannot locate a report that should have been generated, a failover may have occurred. Reports generated during that period are stored on the other appliance. To view those reports, switch to the other appliance.

To access this part of the web UI, your administrator's account access profile must have *Read-Write* permission to items in the *Log & Report* category. For details, see "Restricting permissions" on page 174.

Attack Graphs dashboard

The *Attack Graphs* dashboard contains widgets that display current information about FortiDDoS attack mitigation activity.

These dashboard widgets use the same categories as DDoS Attack Activity reports. For descriptions of the categories, see "DDoS Attack Activity report types" on page 255.

For information on managing the dashboard widgets, see "The dashboard" on page 195.

Diagnostics

The FortiDDoS web UI can display TCP session and packet source diagnostics information.

For both types of information, you can click a column heading to sort the list by the values in that column and drag columns to change their order.

For information on the filter settings for the list, see “Filtering log messages” on page 240.

TCP session statistics

Whenever a service protection profile has one or more active TCP sessions, you can go to *Log & Report* > *Diagnostics* > *Session* to view the following session statistics:

- Source IP Address
- Destination IP Address
- Destination Port
- TCP State
- Connections Count

FortiDDoS displays information for the time indicated by *Generated on*, which is the time you accessed the *Session Diagnostics* page or last clicked *Refresh*.

Use the *Group By* options to select which columns are displayed.

Source statistics

FortiDDoS maintains information about the sources of packets that flow through it, including packets it has blocked, denied, and allowed. It also tracks the IP addresses that are acting as proxies. To view the information, go to *Log & Report* > *Diagnostics* > *Sources*.

To display information about proxy IP addresses only, select *Proxy IP*, and then click *Submit*.

To select the type of information that is displayed in the list, select one or more *Apply flag* options, and then click *Submit*.

Allowed	Displays source information for any packets that the global or service protection profile (SPP) ACL configuration allowed.
Blocked	Displays source information for any packet that the service protection profile (SPP) threshold configuration blocked.
Denied	Displays source information for any packets that the global or service protection profile (SPP) ACL configuration denied.
None	Displays source information for any packets that the appliance configuration did not allow, block, or deny.

FortiDDoS displays information for the time you accessed the *Session Diagnostics* page or last clicked *Refresh*.

The information that is displayed depends on the characteristics of the current traffic and the flag options that you select. If FortiDDoS has not dropped or processed packets that match the selected *Apply flag* categories, no source information is displayed.

Troubleshooting

This topic provides guidelines to help you resolve issues if your FortiDDoS appliance is not behaving as you expect.

Keep in mind that if you cannot resolve the issue on your own, you can contact Fortinet Technical Support:

https://support.fortinet.com

See also

- Solutions by issue type
- Resetting profile data or the appliance configuration
- Restoring firmware (“clean install”)

Solutions by issue type

Recommended solutions vary by the type of issue.

- Connectivity issues
- Resource issues
- Login issues

Fortinet also provides these resources:

- the Release Notes provided with your firmware
- Technical documentation (references, installation guides, and other documents)
- Knowledge base (technical support articles)
- Forums
- Online campus (tutorials and training materials)

Check within your organization. You can save time and effort during the troubleshooting process by checking if other FortiDDoS administrators experienced a similar problem before.

Connectivity issues

One of your first tests when configuring a new Service Protection Profile should be to determine whether non-attack traffic is flowing to your servers.

See also

- Checking hardware connections
- Data path connectivity
- Management network interface connectivity
- Resource issues
- Login issues

Checking hardware connections

If there is no traffic flowing from the FortiDDoS appliance, it may be a hardware problem.

To check hardware connections

- Ensure the network cables are properly plugged in to the interfaces on the FortiDDoS appliance.
- Ensure there are connection lights for the network cables on the appliance.
- Change the cable if the cable or its connector are damaged or you are unsure about the cable's type or quality.
- Connect the FortiDDoS appliance to different hardware to see if that makes a difference.
- In the web UI, select *System* > *Status* > *Dashboard*. In the *System Status* widget, ensure that the status indicators for the ports that are in use are green (indicating that physical connections are present) or flashing green (indicating that data is flowing). Hover over the indicator for further status information.

If any of these checks solve the problem, it was a hardware connection issue. You should still perform some basic software tests to ensure complete connectivity.

If the hardware connections are correct and the appliance is powered on but you cannot connect using the CLI or web UI, you may be experiencing bootup problems. See "Restoring firmware ("clean install")" on page 267.

Data path connectivity

You can use `ping` and other methods to verify that traffic is flowing from the FortiDDoS appliance to the servers it protects.

Verifying the path between client and server

If you are testing connectivity using a client that is directly connected to the appliance, use `ping` to test the traffic flow. For indirect connections, use `traceroute`.

To verify routes between clients and your servers using ping

1. Try to communicate with the server from the client using the `ping` command. Use the following graphs to detect if the traffic has travelled through the FortiDDoS appliance:
   - *Monitor* > *Port Statistics* > *Packets*
   - *Monitor* > *Port Statistics* > *Bits*
   - *Monitor* > *Specific Graphs* > *Protocols* (protocol 1 or 58)

   If you do not see the expected count of bits or packets, continue to the next step.
2. Use the `ping` command on both the client and the server to verify that a route exists between the two. Test traffic movement in both directions: from the client to the server, and the server to the client. Servers do not need to be able to initiate a connection, but must be able to send reply traffic along a return path.

In networks using features such as asymmetric routing, routing success in one direction does ***not*** guarantee success in the other.

To verify routes between clients and your servers using traceroute

Use the `tracert` or `traceroute` command on both the client and the server (depending on their operating systems) to determine if there is a point of failure along the route.

If the route is broken when it reaches the FortiDDoS appliance, first examine its network interfaces and routes. To display network interface information, enter the CLI command:

```
show system interface
```

Testing data path routes & latency with traceroute

`traceroute` sends ICMP packets to test each hop along the route. It sends three packets to the destination, and then increases the time to live (TTL) setting by one, and sends another three packets to the destination. As the TTL increases, packets go one hop farther along the route until they reach the destination.

Most `traceroute` commands display their maximum hop count — that is, the maximum number of steps it will take before declaring the destination unreachable — before they start tracing the route. The TTL setting may result in routers or firewalls along the route timing out due to high latency.

Where `ping` only tells you if the signal reached its destination and returned successfully, `traceroute` shows each step of its journey to its destination and how long each step takes. If you specify the destination using a domain name, the `traceroute` output can also indicate DNS problems, such as an inability to connect to a DNS server.

By default, `traceroute` uses UDP with destination ports numbered from 33434 to 33534. The traceroute utility usually has an option to specify use of ICMP `ECHO_REQUEST` (type 8) instead, as used by the Windows `tracert` utility. If you have a firewall and you want `traceroute` to work from both machines (Unix-like systems and Windows) you will need to allow ***both*** protocols inbound through your firewall (UDP ports 33434 - 33534 and ICMP type 8).

To trace the route to a server from a Microsoft Windows computer

1. Click the *Start* (Windows logo) menu to open it.

   If the host is running Windows XP, instead, go to *Start* > *Run...*
2. Type `cmd` then press Enter.

   The Windows command line appears.

3. Enter the command:

```
tracert {<destination_ipv4> | <destination_fqdn>}
```

If the appliance ***has*** a complete route to the destination, output similar to the following appears:

```
Tracing route to www.fortinet.com [66.171.121.34]
over a maximum of 30 hops:

  1    <1 ms    <1 ms    <1 ms  172.16.1.2
  2     2 ms     2 ms     2 ms  static-209-87-254-221.storm.ca
[209.87.254.221]

  3     2 ms     2 ms    22 ms  core-2-g0-1-1104.storm.ca
[209.87.239.129]
  4     3 ms     3 ms     2 ms  67.69.228.161
  5     3 ms     2 ms     3 ms  core2-ottawa23_POS13-1-0.net.bell.ca
[64.230.164
.17]
(Output abbreviated.)
 15    97 ms    97 ms    97 ms  gar2.sj2ca.ip.att.net [12.122.110.105]
 16     94 ms    94 ms    94 ms  12.116.52.42
 17     87 ms    87 ms    87 ms  203.78.181.10
 18     89 ms    89 ms    90 ms  203.78.181.130
 19     89 ms    89 ms    90 ms  fortinet.com [66.171.121.34]
 20     90 ms    90 ms    91 ms  fortinet.com [66.171.121.34]

Trace complete.
```

Each line lists the routing hop number, the 3 response times from that hop, and the IP address and FQDN (if any) of that hop. Typically a value of `<1ms` indicates a local router.

If the appliance ***does not*** have a complete route to the destination, output similar to the following appears:

```
Tracing route to 10.0.0.1 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  172.16.1.2
  2    <1 ms    <1 ms    <1 ms  172.16.1.10
  3     *        *        *     Request timed out.
  4     *        *        *     Request timed out.
  5 ^C
```

The asterisks (*) and “Request timed out.” indicate no response from that hop in the network routing.

To trace the route to a server from a Linux or Mac OS X computer

1. Open a command prompt.

Alternatively, on Mac OS X, you can use the Network Utility application.

2. Enter (the path to the executable varies by distribution):

```
traceroute {<destination_ipv4> | <destination_fqdn>}
```

If the appliance ***has*** a complete route to the destination, output similar to the following appears:

```
traceroute to www.fortinet.com (66.171.121.34), 30 hops max, 60 byte
packets
 1  172.16.1.2 (172.16.1.2)  0.189 ms  0.277 ms  0.226 ms
 2  static-209-87-254-221.storm.ca (209.87.254.221)  2.554 ms  2.549
ms  2.503 ms
 3  core-2-g0-1-1104.storm.ca (209.87.239.129)  2.461 ms  2.516 ms
2.417 ms
 4  67.69.228.161 (67.69.228.161)  3.041 ms  3.007 ms  2.966 ms
 5  core2-ottawa23_POS13-1-0.net.bell.ca (64.230.164.17)  3.004 ms
2.998 ms  2.963 ms
(Output abbreviated.)
16  12.116.52.42 (12.116.52.42)  94.379 ms  94.114 ms  94.162 ms
17  203.78.181.10 (203.78.181.10)  122.879 ms  120.690 ms  119.049 ms
18  203.78.181.130 (203.78.181.130)  89.705 ms  89.411 ms  89.591 ms
19  fortinet.com (66.171.121.34)  89.717 ms  89.584 ms  89.568 ms
```

Each line lists the routing hop number, the IP address and FQDN (if any) of that hop, and the 3 response times from that hop. Typically a value of `<1ms` indicates a local router.

If the appliance ***does not*** have a complete route to the destination, output similar to the following appears:

```
traceroute to 10.0.0.1 (10.0.0.1), 30 hops max, 60 byte packets
 1  * * *
 2  172.16.1.10 (172.16.1.10)  4.160 ms  4.169 ms  4.144 ms
 3  * * *
 4  * * *^C
```

The asterisks (*) indicate no response from that hop in the network routing.

Relatedly, if the computer's DNS query cannot resolve the host name, output similar to the following appears:

```
example.lab: Name or service not known
Cannot handle "host" cmdline arg `example.lab' on position 1 (argc 1)
```

Management network interface connectivity

Checking routing

`ping` and `traceroute` are useful tools in network connectivity and route troubleshooting for management network interfaces.

Since you typically use these tools for troubleshooting only, allow ICMP (the protocol used by these tools) on interfaces only when you need them. Otherwise, disable ICMP for improved security and performance.

By default, FortiDDoS appliances respond to `ping` and `traceroute`. However, if the appliance does not respond, and there are no firewall policies that block it, ICMP type 0 (`ECHO_REPSPONSE`) might be effectively disabled.

To enable ping and traceroute responses from FortiDDoS

1. Go to *System* > *Network* > *Interface.*

   To access this part of the web UI, you must have *Read-Write* permission in your administrator's account access profile to items in the *System* category. For details, see "Permissions" on page 44.

2. In the row for the management network interface which you want to respond to ICMP type 8 (`ECHO_REQUEST`) for `ping` and UDP for `traceroute`, click *Edit*.

   A dialog appears.

3. Enable *PING*.

Disabling *PING* only prevents FortiDDoS from ***receiving*** ICMP type 8 (`ECHO_REQUEST`) and traceroute-related UDP.

It does ***not*** disable FortiDDoS CLI commands such as `execute ping` or `execute traceroute` that ***send*** such traffic.

4. Click *OK*.

   The appliance should now respond when another device such as your management computer sends a `ping` or `traceroute` to that management interface.

Examining the routing table

When a route does not exist, or when hops have high latency, examine the routing table. The routing table is where the FortiDDoS appliance caches recently used routes.

If a route is cached in the routing table, it saves time and resources that would otherwise be required for a route lookup. If the routing table is full and a new route must be added, the oldest, least-used route is deleted to make room.

To check the routing table for the management network interface in the CLI, enter:

```
diagnose netlink route list
```

Resource issues

If the system resource usage appears to be abnormally high according to the *System Resources widget* or the CLI command:

```
get system status
```

you can view the current consumption by each process by entering this CLI command:

```
diagnose system top 10
```

The above command generates a list of processes every 10 seconds. It includes the process names, their process ID (pid), status, CPU usage, and memory usage.

The report continues to refresh and display in the CLI until you press `q` (quit).

Once you locate an offending PID, you can terminate it:

```
diagnose system kill 9 <pid_int>
```

If the issue recurs, and corresponds with a hardware or configuration change, you may need to change the configuration (especially reducing frequent logging), reduce average traffic load or contact Fortinet Technical Support to prevent the issue from recurring.

See also

- Connectivity issues
- Login issues

Login issues

If the person cannot access the login page at all, it is usually actually a connectivity issue (see "Connectivity issues" on page 260 and "Configuring network interfaces, gateway, and DNS" on page 98) ***unless*** all accounts are configured to accept logins only from specific IP addresses (see "Trusted Host" on page 173).

If the person has lost or forgotten his or her password, the `admin` account can reset other accounts' passwords (see "Changing an administrator's password" on page 175).

When an administrator account cannot log in from a specific IP

If an administrator is entering his or her correct account name and password, but cannot log in from some or all computers, examine that account's trusted host definitions (see "Trusted Host" on page 173). It should include all locations where that person is allowed to log in, such as your office, but should ***not*** be too broad.

See also

- Connectivity issues
- Resource issues

Resetting profile data or the appliance configuration

The following situations are examples of situations that can require a full or partial reset of your appliance configuration:

- The characteristics of the traffic protected by an Service Protection Profile change significantly (for example, you change which server or protocol that it protects)
- You are selling your FortiDDoS appliance
- You are not sure what part of your configuration is causing a problem

You can reset the appliance to its default settings or erase SPP and log data. (If you have not updated the firmware, this is the same as resetting to the factory default settings.)

You can also reset just the thresholds for a profile or the thresholds for a specific OSI layer. For more information, see "Adjusting multiple thresholds at one time" on page 152.

To reset SPP data

Do not shut down the appliance while it is resetting.

1. Go to *Protection Profiles* > *Factory Reset* > *Factory Reset*.
2. For *Service Protection Profile*, select the profile that you want to reset.

3. Select *Reset*.
4. Click *Save*.

   FortiDDoS initializes a whole year's worth of traffic history, even if the appliance has not had a year's worth of uptime. When you erase an SPP's history, its traffic graphs are initialized and do not show past data.
5. Switch the SPP from prevention to detection until FortiDDoS has gathered enough new data to estimate appropriate thresholds for your network.

To reset the appliance's configuration via the web UI

Back up your configuration before beginning this procedure, if possible. For information on backups, see "Backups" on page 167.

Resetting the configuration may include resetting the IP addresses of the network interface that is used for connections to the web UI and CLI. For information on reconnecting to a FortiDDoS appliance whose network interface was reset, see "Connecting to the web UI or CLI" on page 74.

1. Go to *System* > *Status* > *Dashboard*.
2. In the *System Information* widget, click *Reset*.
3. Click *Yes* and wait until the process is complete.

To reset the appliance's configuration via the CLI

Back up your configuration before beginning this procedure, if possible. For information on backups, see "Backups" on page 167.

Resetting the configuration may include resetting the IP addresses of the network interface that is used for connections to the web UI and CLI. For information on reconnecting to a FortiDDoS appliance whose network interface was reset, see "Connecting to the web UI or CLI" on page 74.

Enter the following command:

```
execute factoryreset
```

Alternatively, you can reset the appliance's configuration to its default values for a specific software version by restoring the firmware during a reboot (a "clean install"). See "Restoring firmware ("clean install")" on page 267.

Restoring firmware ("clean install")

Restoring (also called re-imaging) the firmware can be useful in the following cases:

- You are unable to connect to the FortiDDoS appliance using the web UI or the CLI
- You want to install firmware ***without*** preserving any existing configuration (that is, perform a "***clean install***")

Unlike updating firmware, restoring firmware re-images the boot device. Also, restoring firmware can only be done during a boot interrupt, before network connectivity is available, and

therefore ***requires a local console connection to the CLI. It cannot be done through an SSH or Telnet connection.***

Alternatively, if you cannot physically access the appliance's local console connection, connect the appliance's local console port to a terminal server to which you have network access. Once you have used a client to connect to the terminal server over the network, you will be able to use the appliance's local console through it. However, be aware that from a remote location, you may not be able to power cycle the appliance if abnormalities occur.

To restore the firmware

Back up your configuration before beginning this procedure, if possible. Restoring firmware resets the configuration, including the IP addresses of network interfaces. For information on backups, see "Backups" on page 167. For information on reconnecting to a FortiDDoS appliance whose network interface configuration was reset, see "Connecting to the web UI or CLI" on page 74.

1. Download the firmware file from the Fortinet Technical Support web site:

   https://support.fortinet.com/
2. Connect your management computer to the FortiDDoS console port using a RJ-45-to-DB-9 serial cable or a null-modem cable.
3. Initiate a ***local console connection*** from your management computer to the CLI of the FortiDDoS appliance, and log in as the `admin` administrator.

   For details, see "Connecting to the web UI or CLI" on page 74.
4. Connect the MGMT1 port of the FortiDDoS appliance directly or to the same subnet as a TFTP server.
5. Copy the new firmware image file to the root directory of the TFTP server.
6. If necessary, start your TFTP server. (If you do not have one, you can temporarily install and run one such as `tftpd` (Windows, Mac OS X, or Linux) on your management computer.)

Because TFTP is ***not*** secure, and because it does not support authentication and could allow anyone to have read and write access, you should ***only*** run it on trusted administrator-only networks, ***never*** on computers directly connected to the Internet. If possible, immediately turn off `tftpd` off when you are done.

7. Verify that the TFTP server is currently running, and that the FortiDDoS appliance can reach the TFTP server.

   To use the FortiDDoS CLI to verify connectivity, enter the following command:

   ```
   execute ping 192.168.1.168
   ```

   where `192.168.1.168` is the IP address of the TFTP server.
8. Enter the following command to restart the FortiDDoS appliance:

   ```
   execute reboot
   ```
9. As the FortiDDoS appliances starts, a series of system startup messages appear.

   ```
   Press any key to display configuration menu.....
   ```

10. Immediately press a key to interrupt the system startup.

You have only 3 seconds to press a key. If you do not press a key soon enough, the FortiDDoS appliance reboots and you must log in and repeat the `execute reboot` command.

If you successfully interrupt the startup process, the following messages appears:

```
[G]:  Get firmware image from TFTP server.
[F]:  Format boot device.
[B]:  Boot with backup firmware and set as default.
[Q]:  Quit menu and continue to boot with default firmware.
[H]:  Display this list of options.

Enter G,F,B,Q,or H:

Please connect TFTP server to Ethernet port "1".
```

11. If the firmware version requires that you first format the boot device before installing firmware, type `F`. Format the boot disk before continuing.
12. Type `G` to get the firmware image from the TFTP server.

    The following message appears:

    ```
    Enter TFTP server address [192.168.1.168]:
    ```

13. Type the IP address of the TFTP server and press Enter.

    The following message appears:

    ```
    Enter local address [192.168.1.188]:
    ```

14. Type a temporary IP address that can be used by the FortiDDoS appliance to connect to the TFTP server.

    The following message appears:

    ```
    Enter firmware image file name [image.out]:
    ```

15. Type the file name of the firmware image and press Enter.

    The FortiDDoS appliance downloads the firmware image file from the TFTP server and displays a message similar to the following:

    ```
    MAC:00219B8F0D94
    ############################
    Total 28385179 bytes data downloaded.
    Verifying the integrity of the firmware image..
    Save as Default firmware/Backup firmware/Run image without
    saving:[D/B/R]?
    ```

If the download fails after the integrity check with the error message:

```
invalid compressed format (err=1)
```

but the firmware matches the integrity checksum on the Fortinet Technical Support web site, try a different TFTP server.

16. Type `D`.

    The FortiDDoS appliance downloads the firmware image file from the TFTP server. The FortiDDoS appliance installs the firmware and restarts. The time required varies by the size of the file and the speed of your network connection.

    The FortiDDoS appliance reverts the configuration to default values for that version of the firmware.

17. To verify that the firmware was successfully installed, log in to the CLI and type:

    ```
    get system status
    ```

    The firmware version number is displayed.

18. Either reconfigure the FortiDDoS appliance or restore the configuration file. For details, see "How to set up your FortiDDoS" on page 53 and "Restoring a previous configuration" on page 169.

If you are ***downgrading*** the firmware to a previous version, and the settings are not fully backwards compatible, the FortiDDoS appliance may either remove incompatible settings, or use the feature's default values for that version of the firmware. You may need to reconfigure some settings.

Installing firmware overwrites any FortiGuard IP Reputation Service definitions and disables the service. After any firmware update, re-enable the IP Reputation feature. FortiDDoS downloads current definitions as part of the enabling process. For more information, see "FortiGuard IP Reputation Service" on page 130.

Appendix A: Port numbers

Communications between the FortiDDoS appliance, clients, servers, and FortiGuard Distribution Network (FDN) require that any routers and firewalls between them permit specific protocols and port numbers.

The following tables list the default port assignments used by FortiDDoS. Many may differ if you have changed them. For example, to change the port numbers used by the web UI and CLI, see "Global web UI & CLI settings" on page 46.

Table 15: Default ports used by FortiDDoS for outgoing traffic

Port Number	IP Protocol Number/ Service	Purpose
N/A	ARP	HA failover of network interfaces. See "Heartbeat link and synchronization" on page 65.
25	TCP	SMTP for alert email. See "Alert email" on page 241.
53	UDP	DNS queries. See "Configuring DNS settings" on page 105.
69	UDP	TFTP for backups, restoration, and firmware updates. See commands such as `execute backup` or `execute restore`.
123	UDP	NTP synchronization. See "Setting the system time & date" on page 95.
162	UDP	SNMP traps. See "SNMP traps & queries" on page 243.
443	TCP	FortiGuard polling and update downloads. See "FortiGuard IP Reputation Service" on page 130.
514	UDP	Syslog. See "Configuring logging to a remote logging server" on page 237.
6055	UDP	HA heartbeat. multicast. See "Heartbeat link and synchronization" on page 65.
6056	UDP	HA configuration synchronization. multicast. See "Heartbeat link and synchronization" on page 65.

Table 16: Default ports used by FortiDDoS for incoming traffic (listening)

Port Number	IP Protocol Number/ Service	Purpose
N/A	ICMP	`ping` and `traceroute` responses. See "Configuring the network interfaces" on page 98.
22	TCP	SSH administrative CLI access. See "Configuring the network interfaces" on page 98.

Table 16: Default ports used by FortiDDoS for incoming traffic (listening)

Port Number	IP Protocol Number/ Service	Purpose
23	TCP	Telnet administrative CLI access. See “Configuring the network interfaces” on page 98.
80	TCP	HTTP administrative web UI access. See “Configuring the network interfaces” on page 98 and “How to use the web UI” on page 43.
161	UDP	SNMP queries. See “SNMP traps & queries” on page 243 and “Configuring the network interfaces” on page 98.
443	TCP	HTTPS administrative web UI access. Only occurs if the destination address is a network interface’s IP address. See “Configuring the network interfaces” on page 98 and “How to use the web UI” on page 43.
6055	UDP	HA heartbeat. multicast. See “Heartbeat link and synchronization” on page 65.
6056	UDP	HA configuration synchronization. multicast. See “Heartbeat link and synchronization” on page 65.

Appendix B: Switch & router configuration

Switch configuration for load balancing

The following example load balancing configuration is for the FortiSwitch 248-B DPS Ethernet switch.

It configures two trunk groups with eight ports per trunk. Trunk 10 is used for Internet traffic and trunk 11 is used for server-side traffic.

You use the `load-balance-hash` command to specify `src-dst-ip-ipports` as the hash distribution algorithm (hash mode) to apply to all trunk groups. This mode uses a 4-tuple (source and destination IP address and source IP L4 port and destination IP L4 port) to ensure that all packets belonging to a session pass through the same port pair on FortiDDoS appliance in both directions.

For more information about load balancing with FortiDDoS, see See "Load balancing" on page 58.

```
(clientSide-84.82) #show run
!Current Configuration:
!
!System Description "FortiSwitch-248B-DPS 48x1G & 4x10G"
!System Software Version "5.2.0.2.4"

serviceport ip 192.168.22.98 255.255.255.0 0.0.0.0
vlan database
vlan name 10 "egress"
vlan name 11 "ingress"
exit

port-channel "egress" 1
interface 0/1
channel-group 1/1
exit
interface 0/3
channel-group 1/1
exit
interface 0/5
channel-group 1/1
exit
interface 0/7
channel-group 1/1
exit
interface 0/9
channel-group 1/1
exit
interface 0/11
channel-group 1/1
```

```
exit
interface 0/13
channel-group 1/1
exit
interface 0/15
channel-group 1/1
exit
port-channel "ingress" 2
interface 0/2
channel-group 1/2
exit
interface 0/4
channel-group 1/2
exit
interface 0/6
channel-group 1/2
exit
interface 0/8
channel-group 1/2
exit
interface 0/10
channel-group 1/2
exit
interface 0/12
channel-group 1/2
exit
interface 0/14
channel-group 1/2
exit
interface 0/16
channel-group 1/2
exit

mac-addr-table aging-time 60000

interface 0/1
no cdp run
switchport allowed vlan add 10
exit

interface 0/2
no cdp run
exit
interface 0/3
no cdp run
exit

interface 0/4
no cdp run
```

```
exit

interface 0/5
no cdp run
exit

interface 0/6
no cdp run
exit

interface 0/7
no cdp run
exit

interface 0/8
no cdp run
exit

interface 0/9
no cdp run
exit

interface 0/10
no cdp run
exit

interface 0/11
no cdp run
exit

interface 0/12
no cdp run
exit

interface 0/13
no cdp run
exit

interface 0/14
no cdp run
exit

interface 0/15
no cdp run
exit

interface 0/16
no cdp run
exit
```

```
interface 0/17
no cdp run
switchport allowed vlan add 10
switchport native vlan 10
exit

interface 0/18
no cdp run
switchport allowed vlan add 11
switchport native vlan 11
exit

interface 0/49
no cdp run
switchport allowed vlan add 10
switchport native vlan 10
exit

interface 0/50
no cdp run
switchport allowed vlan add 11
switchport native vlan 11
exit

interface 1/1
staticcapability
switchport allowed vlan add 10
switchport native vlan 10
lacp collector max-delay 0
exit

interface 1/2
staticcapability
switchport allowed vlan add 11
switchport native vlan 11
lacp collector max-delay 0
exit

interface 1/3
staticcapability
switchport allowed vlan add 10
switchport tagging 10
lacp collector max-delay 0
exit

interface 1/4
staticcapability
switchport allowed vlan add 11
```

```
switchport tagging 11
lacp collector max-delay 0
exit

router rip
exit
router ospf
exit
exit

(clientSide-84.82) #
(clientSide-84.82) #show load-balance
Hash Mode: src-dst-ip-ipport
```

Configuring the routers & switch for traffic diversion

The following example router and switch configuration is used for traffic diversion.

For more information about traffic diversion for FortiDDoS, see “Traffic diversion” on page 60.

Router configuration

```
vlan internal allocation policy ascending
!
interface GigabitEthernet1/0/1
  no switchport
  no ip address
!
interface GigabitEthernet1/0/2
  switchport access vlan 3
  switchport trunk encapsulation dot1q
!
interface GigabitEthernet1/0/3
!
interface GigabitEthernet1/0/4
!
interface GigabitEthernet1/0/5
!
interface GigabitEthernet1/0/6
!
interface GigabitEthernet1/0/7
!
interface GigabitEthernet1/0/8
!
interface GigabitEthernet1/0/9
!
interface GigabitEthernet1/0/10
  switchport access vlan 2
```

```
!
interface GigabitEthernet1/0/11
ip address 10.100.0.250 255.255.255.0
no ip directed-broadcast
ip policy route-map FDD-X00A-PBR
!
interface GigabitEthernet1/0/12
!
interface Vlan2
  ip address 10.1.0.251 255.255.255.0
!
interface Vlan3
  ip address 192.168.100.51 255.255.255.0
!
!
ip classless
ip route 207.117.1.0 255.255.255.0 10.1.0.250
!
!
ip access-list extended zone-A
permit ip any  207.117.0.0 0.0.0.255
!
route-map FDD-X00A-PBR permit 100
match ip address zone-A
set ip next-hop 10.200.0.254
!
route-map FDD-X00A-PBR permit 101
description let thru all other packets without modifying next-hop
```

Switch configuration

```
interface GigabitEthernet1/0/1
  no switchport
  no ip address
  channel-group 1 mode on
!
interface GigabitEthernet1/0/2
  switchport access vlan 3
  switchport trunk encapsulation dot1q
!
interface GigabitEthernet1/0/3
  switchport access vlan 3
!
interface GigabitEthernet1/0/4
  switchport access vlan 3
  switchport trunk encapsulation dot1q
!
interface GigabitEthernet1/0/5
```

```
!
interface GigabitEthernet1/0/6
!
interface GigabitEthernet1/0/7
!
interface GigabitEthernet1/0/8
!
interface GigabitEthernet1/0/9
!
interface GigabitEthernet1/0/10
  switchport access vlan 2
!
interface GigabitEthernet1/0/11
  switchport access vlan 5
!
interface GigabitEthernet1/0/12
  switchport access vlan 4
!
!

interface Vlan1
  no ip address
!
interface Vlan3
  ip address 192.168.100.50 255.255.255.0
!
interface Vlan4
  ip address 10.100.0.250 255.255.255.0
!
!
interface Vlan5
  ip address 10.1.0.250 255.255.255.0
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.100.0.254
```

Index

C

D

E

F

G

H

I

J

K

Q

R

S

T

U

V

W

X

Z